


Volcano, an Extensible and Parallel Query Evaluation System

Goetz Graefe

CU-CS-481-90



University of Colorado at Boulder

DEPARTMENT OF COMPUTER SCIENCE

ANY OPINIONS, FINDINGS, AND CONCLUSIONS OR RECOMMENDATIONS EXPRESSED IN THIS PUBLICATION ARE THOSE OF THE AUTHOR(S) AND DO NOT NECESSARILY REFLECT THE VIEWS OF THE AGENCIES NAMED IN THE ACKNOWLEDGMENTS SECTION.

**Volcano, an Extensible and Parallel
Query Evaluation System**

Goetz Graefe

CU-CS-481-90 July 1990

Department of Computer Science
Campus Box 430
University of Colorado at Boulder
Boulder, Colorado 80309-430
Telephone: 303-492-7514
FAX: 303-492-2844

Volcano, an Extensible and Parallel Query Evaluation System

Goetz Graefe
University of Colorado at Boulder

Abstract

To investigate the interactions of extensibility and parallelism in database query processing, we have developed a new dataflow query execution system called Volcano. The Volcano effort also provides an environment for research and education in database systems design, heuristics for query optimization, and resource allocation. At the current time, Volcano is operational on a variety of hardware platforms, currently on single-processor and shared-memory machines, and soon also on distributed-memory parallel machines.

Following a design principle well-established in operating systems research but not exploited in most database system designs, Volcano provides *mechanisms* to support *policies*. Policies can be set by a human experimenter or by a query optimizer. The separation of mechanisms and policies has contributed to the extensibility and modularity of modern operating systems, and may make the same contribution to extensible database systems.

Volcano uses a standard interface between algebra operators, allowing easy addition of new operators and operator implementations. Operations on individual items, e.g., predicates, are imported into the query processing operators using *support functions*. The interface to support functions allows both compiled and interpreted query execution. The semantics of support functions is not prescribed; any data type including complex objects and any operation can be realized. Thus, Volcano is *extensible* with new operators, algorithms, data types, and type-specific methods.

For efficiency in single-process query execution, transfer of data and control between operators is based on the *iterator* paradigm, also called lazy evaluation, demand-driven dataflow, or synchronous pipelines. In addition to an operator set comparable to the query processing engines of most database systems, Volcano includes two novel *meta-operators*. Meta-operators do not perform semantic data transformations but provide query processing control. The *choose-plan* meta-operator supports *dynamic query evaluation plans* that allow delaying selected optimization decisions until run-time, e.g., for embedded queries with free variables. The *exchange* meta-operator supports *intra-operator parallelism* on partitioned datasets and both *vertical and horizontal inter-operator parallelism*, translating between demand-driven dataflow within processes and data-driven dataflow between processes.

All operators, with the exception of the exchange operator, have been designed, implemented, debugged, and tuned in a single-process environment, and parallelized using the exchange operator. Even operators not yet designed can be parallelized using this new operator if they use and provide the iterator interface. Thus, the issues of data manipulation and parallelism have become *orthogonal*, making Volcano the first implemented query execution engine that effectively combines extensibility and parallelism.

1. Introduction

In order to investigate the interactions of extensibility, efficiency, and parallelism in database query processing, and to provide a testbed for database systems research and education, we have designed and implemented a new query evaluation system called Volcano. It is intended to provide an experimental vehicle for research into query execution techniques and query optimization heuristics rather than a database system ready to support applications. It is not a complete database system as it lacks features such as a user-friendly query language, a type system for instances (record definitions), a query optimizer, or catalogs. Because of this focus, Volcano is able to serve as experimental vehicle for a multitude of purposes, all of them open-ended, which results in a combination of require-

ments that have not been integrated in a single system before. First, it is modular and extensible to enable future research, e.g., on algorithms, data models, resource allocation, parallel execution, load balancing, and query optimization heuristics. Thus, Volcano provides an infrastructure for experimental research rather than a final research prototype in itself. Second, it is simple in its design to allow student use and research. Modularity and simplicity are very important for this purpose because they allow student working on projects without understanding of the entire design and all its details, as well as several concurrent student projects. Third, Volcano's design does not presume any particular data model; the only assumption is that query processing is based on transforming sets of items using parameterized operators. To achieve data model independence, the design very consistently separates set processing *control* (which is provided and inherent in the Volcano operators) from *interpretation and manipulation* of data items (which is imported into the operators, as described later). Fourth, to free algorithm development, tuning, and initial experimentation from the intricacies of parallelism but also to allow experimentation with parallel query processing, Volcano can be used as a single-process or as a parallel system. Single-process query evaluation plans can already be parallelized easily on shared-memory machines and soon also on distributed-memory machines. Fifth, Volcano is realistic in its query execution paradigm to ensure that students learn how query processing is really done in commercial database products. For example, using temporary files to transfer data from one operation to the next as suggested in most text books has a substantial performance penalty, and is therefore used neither in real database systems nor in Volcano. Finally, Volcano's means for parallel query processing could not be based on existing models since all methods explored to-date have been designed with a particular data model and operator set in mind. Instead, our design goal was to make parallelism and data manipulation *orthogonal*, which means that the mechanisms for parallel query processing are independent of the operator set and semantics, and that all operators, including new ones, could be designed and implemented independently of future parallel execution.

Since its very purpose is to allow future extensions and research, Volcano is continuously being modified and extended¹. Among the most important recent extensions were the design and implementation of two *meta-operators*. Both of them are not only new operators but also embody and encapsulate new concepts for query processing. They are meta-operators since they don't contribute to data manipulation, selection, derivation, etc., but

¹ In fact, it is currently being used in five academic institutions, three industrial firms, and one non-profit organization.

instead provide additional control over query processing that cannot be provided by conventional operators. The *choose-plan* operator implements *dynamic query evaluation plans*, a concept developed for queries that must be optimized with incomplete information [50]. For example, it is not possible to reliably optimize an embedded query if one of the constants in the query predicate is actually a program variable and therefore unknown during compilation and optimization. Dynamic plans allow preparation of multiple equivalent plans, each one optimal for a certain range of actual parameter values. The *choose-plan* operator selects among these plans at run-time while all other operators in Volcano's operator set (present or future) are entirely oblivious to the presence and function of the *choose-plan* operator. The second meta-operator, the *exchange* operator, implements and controls parallel query evaluation in Volcano. While operators can exchange data without the *exchange* operator, in fact within processes as easily as a single procedure call, this new operator exchanges data across process and processor boundaries. All other operators are implemented and execute without regard to parallelism; all parallelism issues like partitioning and flow control are encapsulated in and provided by the *exchange* operator. Thus, data manipulation and parallelism are indeed orthogonal in Volcano [53]. Beyond the cleanliness from a software engineering point of view, it was also very encouraging to see that this method of parallelizing a query processing engine did indeed allow linear or near-linear speedup [55].

This article is a general overview describing the overall goals and design principles. Other articles on Volcano were written on special aspects of the system, e.g. [48-50, 53, 55, 56, 69]. These articles also include experimental performance evaluations of Volcano's techniques and algorithms, in particular [55].

The present article is organized as follows. In the following section, we briefly review previous work that influenced Volcano's design. A detailed description of Volcano follows in Section 3. Section 4 contains a discussion of extensibility in the system. Dynamic query evaluation plans and their implementation are described in Section 5. Parallel processing encapsulated in the *exchange* module is described in Section 6. Section 7 contains a summary and our conclusions from this effort.

2. Related Work

Since so many different systems have been developed to process large datasets efficiently, we only survey the systems that have significantly influenced the design of Volcano. Our work has been influenced most strongly by WiSS, GAMMA, and EXODUS. The Wisconsin Storage System (WiSS) [27] is a record-oriented file system providing heap files, B-tree and hash indices, buffering, and scans with predicates. GAMMA [31, 34] is a software

database machine running on a number of general-purpose computers as a backend to a UNIX host machine. It was developed on 17 VAX 11/750's connected with each other and the VAX 11/750 host via a 80 Mb/s token ring. Eight GAMMA processors had a local disk device, accessed using WiSS. The disks were accessible only locally, and update and selection operators used only these eight processors. The other, diskless processors were used for join processing [41]. Recently, the GAMMA software has been ported to an Intel iPSC/2 hypercube with 32 nodes, each with a local disk drive [34]. GAMMA uses hash-based algorithms extensively, implemented in such a way that each operator is executed on several (usually all) processors and the input stream for each operator is partitioned into disjoint sets according to a hash function [30, 31, 41]. A detailed description of GAMMA can be found in [31, 41], more information on its performance in [32, 34, 42].

The limited data model and extensibility of the GAMMA machine led to the search for a more flexible but equally powerful query processing model. The operator design used in the GAMMA database machine software gives each operator control within its own process, leaving it to the networking and operating system software to synchronize multiple operators in producer-consumer relationships using flow-control mechanisms. This design, while working extremely well in GAMMA, does not lend itself to single-process query evaluation since multiple loci of control, i.e., multiple operators, cannot be realized inside a single process without special pseudo-multi-process mechanisms such as threads. Therefore, GAMMA's operator and data transfer concepts are not suitable for an efficient query processing engine intended for both sequential and parallel query execution.

EXODUS [18, 20, 22] is an extensible database system with some components following the "tool-kit" approach, e.g., the optimizer generator [44-46] and the E database implementation language [98-100], and other components built as powerful but fixed components, e.g., the storage manager [19]. Originally, EXODUS was conceived to be data-model-independent, i.e., it was supposed to support a wide variety of data models, but later a novel, powerful structurally object-oriented data model called Extra was developed [21, 130]. The concept of data model independence as first explored in EXODUS has been retained in the Volcano project and the design and implementation of its software. During the design of the EXODUS storage manager, many storage and access issues explored in WiSS and GAMMA were revisited. Lessons learned and tradeoffs explored in these discussions certainly helped in forming the ideas behind Volcano. The design and development of E influenced the strong emphasis on iterators for query processing. Part of the original motivation for developing Volcano was the desire for a data-model-independent query execution engine as complement to the EXODUS optimizer generator.

A number of further conventional (relational) and extensible systems have influenced our design. Ingres [114, 115] and System R [2, 24] have probably influenced most database systems, in particular their extensible follow-on projects Starburst [61, 105] and Postgres [116, 119]. It is interesting to note that independently of our work the Starburst group has also identified the demand-driven iterator paradigm as a suitable basis for an extensible single-process query evaluation architecture [60] after using it successfully in the System R relational system. GENESIS [5, 6] early on stressed the importance of uniform operator interfaces for extensibility and software reusability.

XPRS [117, 118] was the first project aiming to combine extensibility and parallelism. Its basic premise is to implement Postgres on top of RAID disk arrays [92] and the Sprite operating system [90, 91]. XPRS and GAMMA differ basically in four ways. First, GAMMA supports a purely relational data model while XPRS supports an extensible relational model, Postgres. Second, GAMMA's main form of parallelism is intra-operator parallelism based on partitioned data sets. XPRS, on the other hand, will rely on bushy parallelism, i.e., concurrent execution of different subtrees in a complex query evaluation plan. Third, GAMMA relies heavily on hashing for joins and aggregations whereas XPRS will have a mainly sort-based query processing engine [118]. Fourth, GAMMA is built on the premise that distributed memory is required to achieve linear speed-up [15] while XPRS is being implemented on a shared-memory machine.

Both XPRS and Volcano combine parallelism and extensibility, but XPRS is a far more comprehensive project than Volcano. In particular, XPRS includes a data model and a query optimizer. On the other hand, Volcano is more extensible precisely because it does not presume a data model. Therefore, Volcano could be used as the query processing engine in a parallel extensible-relational system such as XPRS. Moreover, it will eventually include a data-model-independent optimizer generator to form a complete query processing research environment.

3. Volcano System Design

In this section, we provide an overview of the design of Volcano. At the current time, Volcano is a library of about two dozen modules with a total of about 15,000 lines of C code. These modules include a file system, buffer management, sorting, B⁺-trees, and two algorithms each (sort- and hash-based) for natural join, semi-join, all three outer joins, anti-join, aggregation, duplicate elimination, union, intersection, difference, anti-difference, and relational division. Moreover, two modules implement dynamic query evaluation plans and allow parallel processing of all algorithms listed above.

All operations on individual records are deliberately left open for later definition. Instead of inventing a language in which to specify selection predicates, hash functions, etc., functions are passed to the query processing operators to be called when necessary with the appropriate arguments. These *support functions* are described later in more detail. One common and repeating theme in the design of Volcano is that it provides *mechanisms* for query evaluation to allow selection of and experimentation with *policies*. The separation of mechanisms and policies is a very well-known and well-established principle in the design and implementation of operating systems [76], but it has not been used as extensively and consistently in the implementation of database systems. It has contributed significantly to the extensibility and modularity of modern operating systems, and may make the same contribution to extensible database systems.

Currently, Volcano consists of two layers, the file system layer and the query processing layer. The former provides record, file, and index operations including scans with optional predicates, and buffering; the latter is a set of query processing modules that can be nested to build complex query evaluation trees. This separation can be found in most query evaluation systems, e.g., RSS and RDS in System R [2, 24] and Core and Corona in Starburst [61, 105]. System catalogs or a data dictionary are not included in Volcano since the system was designed to be extensible and independent from any particular data model. We start our description at the bottom, the file system, and then discuss the query processing modules.

3.1. The File System

Within our discussion of the Volcano file system, we also proceed bottom-up, from devices through buffer management to data files and indices. The existing facilities are meant to provide a backbone of a query processing system, and are designed such that modifications and additions can easily be accomplished as the need arises.

Volcano stores data on disk devices. Devices are viewed as arrays of fixed-sized pages. The page size is a compile-time constant, e.g., 1 KB. UNIX files are used to simulate devices, similar to WiSS. It is quite straightforward to port Volcano to using raw devices or other primitive operating system services and has been done successfully for performance evaluations. It is assumed that a contiguous array of pages can be read or written with a single read or write system call.

The *buffer manager* is the most interesting part of the file system. Because buffer management is performance-critical in any database system, the Volcano buffer manager was designed to include mechanisms that can be used most effectively and efficiently in a large variety of contexts and with a wide array of policies. In

consequence, its features include multiple buffer pools, variable-length units of buffering which are called *clusters* in Volcano, and replacement hints from the next higher software level.

Buffer pools can be created and destroyed dynamically, and multiple buffer pools can exist concurrently. When a device is mounted, i.e., created or opened, the device is assigned to a buffer pool. When a pool is created, it is created with a *norm size*, e.g., 512 KB, and an *actual size* of 0. The buffer pool grows with each new request that cannot be satisfied with the resident clusters until the actual size reaches the norm size and page replacement commences. In the case of a buffer fault, it first is checked whether the cluster at the bottom of the LRU stack is of the same size as the requested cluster. If so, it is replaced. If not, or if there is no free cluster in the stack because the buffer is overcommitted, space for the new cluster is allocated. For this purpose, there are more cluster descriptors in the buffer pool than the norm size of the pool, and the buffer pool can grow beyond its norm size. After the requested cluster is loaded into the buffer, the buffer deallocates clusters from the bottom of the LRU stack until the actual size is below or equal to the norm size, or until the LRU stack is empty.

A number of investigations concerned database buffer replacement strategies, as surveyed in [36]. Buffer replacement in systems with variable-size pages was recently investigated by Sikeler [109]. To comment only on one investigation, Chou performed an extensive comparison of replacement algorithms using simulations driven by real database system reference traces [25, 26]. However, a close examination of Chou's work reveals that most of the benefits of Chou's algorithm DBMIN could be realized by using *hints* from higher level software, e.g., the query processing routines, for each cluster when a cluster is pinned or released². Thus, when pinning or unpinning a cluster in the buffer, a hint can be given to Volcano's buffer manager whether to *keep* or to *toss* this cluster. This hint is used by the buffer manager to decide whether to insert the cluster at the top or at the bottom of the LRU replacement list. In effect, the replacement policy is a switched LRU/MRU policy. The LRU policy is implemented in such a way that dirty clusters are paged out slower than clean ones since their replacement cost is higher.

Incidentally, this hint facility is similar to the *LOVE/HATE* hint used in Starburst's buffer manager [61]. However, while the default in Starburst is *LOVE* (LRU), the default in Volcano is *toss* (MRU). The MINIX operating system [122] also inserts pages either at the top or the bottom of a LRU stack for data and indirect pages,

² Chou's algorithm requires a hint when a scan is opened about which of a set of access patterns can be expected. Chou's algorithm also includes load control by detecting an overcommitted buffer and the danger of thrashing. However, we believe that load control can easily be added to Volcano simply by estimating hot set sizes [101] and keeping track of their sum, or, even simpler, by observing and limiting disk utilization [25].

respectively. The decision is static, however, and does not accept hints from the user level.

The buffer manager's hint facility is an excellent example of Volcano's design principle to implement *mechanisms* to support multiple *policies*. The buffer manager only provides the mechanisms, i.e., pinning, page replacement, and reading and writing disk pages, while the higher-level software determines the policies depending on data semantics, importance, and access patterns. It is surprising that database buffer managers derive replacement decisions from *observed* reference behavior in spite of the fact that this behavior is generated by higher-level database software and thus known and foreseeable in advance within the same system, albeit different subcomponents.

Hints to a shared resource manager like a buffer manager always carry two dangers. First, in accepting a hint, the resource manager must *trust* the hinting entity. While this may be a problem for an operating system that has to serve a wide user community, a database buffer manager only serves database modules which are designed and written by the same people as the buffer manager code. Therefore, in the case of a database buffer manager, even in an extensible database system, trust should not be an issue. The second danger is that hints may be generated *locally* while the buffer manager should improve I/O performance *globally*. This can result in a buffer client, e.g., a scan, being too "shy" about asking for more resources in times of low contention, and too "demanding" during high contention. In other words, the local decision which hint to generate is typically oblivious to the global situation. This is the reason why the replacement hints are *hints* rather than *directives* or *absolutes*: the buffer manager is free to ignore any hint if the global situation requires it to do so. Another issue is that local decisions are made without knowledge of the global reference behavior. For example, if each database transaction touches a certain catalog page once, but only once, no transaction will issue a *keep* hint. However, it is clearly worthwhile to keep this page in the buffer. Therefore, the buffer manager must make some decisions independent of the hint facilities, and these decisions do indeed rely on observing reference behavior. In the Volcano buffer manager, a page is automatically switched to *keep* if it is reclaimed from the free list.

Files are composed of records, clusters, and extents. Since file operations are invoked very frequently in any database system, all design decisions in the file module have been made to provide basic functionality with the highest attainable performance. A cluster, consisting of one or more pages, is the unit of I/O and of buffering, as discussed above. The cluster size is fixed for each file individually. Thus, different files on the same device can have different cluster sizes. Records must fit into clusters to avoid the difficulties and inefficiencies of spanning

records, i.e., records that are divided over multiple clusters. Since clusters can be very large, however, this is not a restriction for most applications. Disk space for files is allocated in physically contiguous extents, because extents allow very fast scanning without seeks and large-chunk read-ahead and write-behind.

Records are identified by a *record identifier* (RID), and can be accessed directly using the RID³. There are two modes of record insertion. If a *near-RID* is provided as a hint with the insertion request, the new record is inserted into the old record's cluster if possible. This hint facility allows user-level clustering of records and objects. By default, the record is inserted into an arbitrary cluster of the file. For this purpose the file descriptor includes a hint pointing to a cluster that has some free space. If no such cluster can be found, a new cluster is appended to the file.

For fast access to a large sets of records, Volcano supports not only individual file and record operations but also *scans* that support read-next and append operations. There are two interfaces to file scans; one is part of the file system and is described momentarily; the other is part of the query processing level and is described later. The first one has the standard procedures for file scans, namely *open*, *next*, *close*, and *rewind*. The *next* procedure returns the main memory address of the next record. This address is guaranteed (pinned) until the next operation is invoked on the scan. A file scan supports the notion of a *current record* and *current cluster*, the cluster which contains the current record. The current record is the record returned by the last *next* call. The current cluster is fixed in the buffer while the scan is open. Thus, getting the next record within the same cluster does not require calling the buffer manager and can be performed very efficiently in about 100 instructions.

For fast creation of files, scans support an *append* operation. It allocates a new record slot, either in the current cluster or in a new cluster appended to the end of the file, and returns the new slot's main memory address. It is the caller's responsibility to fill the provided record space with useful information, i.e., the *append* routine is entirely oblivious to the data and their representation. Since *append* can use the current cluster in most cases (without invoking the buffer manager) and since it does not perform any copying, it is about as fast as the *next* operation.

³ A RID contains device, page, and record numbers plus the cluster size. The cluster size allows direct access to the cluster without any additional information, e.g., about an open file.

Scans also support optional *predicates*. Predicates are passed to the scan by means of a function entry point and a typeless pointer which serves as a predicate argument. The predicate function is called by the *next* procedure with the argument and a record address. Selective scans are the first example of *support functions* mentioned briefly in the introduction. Instead of determining a qualification itself, the scan mechanism relies on a function *imported* from a higher level.

Arguments to support functions can be used in two ways, namely in compiled and interpreted query execution. While Katz has shown that compiled query plans result in better execution performance [68], interpretation provides easier portability and allows startup with less overhead. In compiled scans, i.e., when the predicate evaluation function is available in machine code, the argument can be used to pass a constant or a pointer to several constants to the predicate function. For example, if the predicate consists of comparing a record field with a string, the comparison function is passed as predicate function while the search string is passed as predicate argument. In interpreted scans, i.e., when a general interpreter is used to evaluate all predicates in a query, they can be used to pass appropriate code to the interpreter. The interpreter's entry point is given as predicate function. Thus, both interpreted and compiled scans are supported with a single simple and efficient mechanism. Volcano's use of support functions and their arguments is another example for a mechanism that leaves a policy decision, in this case whether to use compiled or interpreted scans, open to be decided by higher-level software.

Indices are implemented currently only in the form of top-down B⁺-trees [58]. Hash indices will be added shortly since they provide superior performance for exact-match lookup. Multi-dimensional indices may be added at a later date. In B⁺-trees, the sizes of internal index nodes and leaf clusters can be different multiples of a page. We expect that larger nodes and leaves will provide better performance for large keys, even without key compression. A leaf entry consists of a key and information. The information part typically is a RID, but it could include more or different information. The key and the information can be of any type; a comparison function must be provided to compare keys. The comparison function uses an argument equivalent to the one described for scan predicates. Permitting any information in the leaves gives more choices in physical database design. It is another example of Volcano providing a mechanism to allow a multitude of designs and usage policies.

B⁺-trees support scans similar to files, including predicates. In addition, B⁺-tree scans allow seeking to a particular key, and setting lower and upper bounds. Finally, B⁺-trees also support an *append* operation for fast loading. Its design assumes that entries are appended in ascending key order. Root-to-leaf traversals are necessary only

when appending a new leaf to the B⁺-tree. The *append* operation includes parameters for free space in B⁺-tree nodes and leaves.

For intermediate results in query processing (later called *streams*), Volcano uses special devices called *virtual devices*. The difference between virtual and disk devices is that data pages of virtual devices only exist in the buffer. As soon as such data pages are unpinned, they disappear and their contents are lost. Thus, Volcano uses the same mechanisms and function calls for permanent and intermediate data sets, easing implementation of new operators significantly.

In summary, much of Volcano's file system is conventional in its goals but implemented in a flexible, efficient, and compact way. The file system supports basic abstractions and operations, namely devices, files, records, B⁺-trees, and scans. It provides mechanisms to access these objects, leaving many policy decisions to higher-level software. High performance was a very important goal in the design and implementation of these mechanisms since future performance studies and parallelization only make sense if the underlying mechanisms are efficient. Furthermore, research into implementation and performance tradeoffs for extensible database systems and new data models is only relevant if an efficient evaluation platform is used.

3.2. Query Processing

The file system routines described above are utilized by the query processing routines to evaluate complex queries. Queries are expressed as *query plans* or algebra expressions; the operators of this algebra are query processing algorithms and we call the algebra an *executable algebra* as opposed to *logical algebras*, e.g., relational algebra. We will describe the operations using relational terminology hoping that this will assist the reader. We must point out, however, that the operations can be viewed and are implemented as operations on sets of objects, and that Volcano does not depend on assumptions about the internal structure of such objects. In fact, we intend to use Volcano for query processing in an object-oriented database system [47, 54]. The key to this use of Volcano is that set processing and interpretation of data items are separated.

At the present time, Volcano does not include mechanisms to create query plans. In a complete database system, plans are created in a sequence of steps illustrated in Figure 1. The output of the optimization step is called *query evaluation plan*, *access plan*, or simply *plan*. It can be expressed in many languages, e.g., in ASL used in System R [77] or as complex algebra expression as in Volcano.

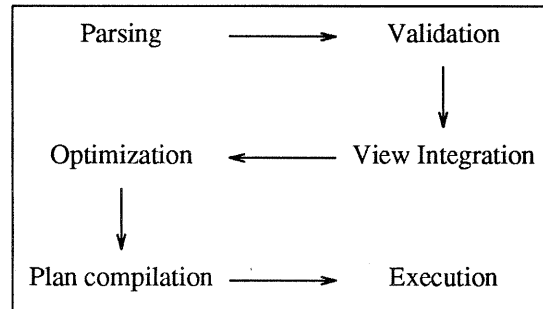


Figure 1. Query Preparation.

In Volcano, all algebra operators are implemented as *iterators*, i.e., they support a simple *open-next-close* protocol. Basically, iterators provide the iteration component of a loop, i.e., initialization, increment, loop termination condition, and final housekeeping. These functions allows "iteration" over the results of any operation similar to the iteration over the result of a conventional file scan. Associated with each iterator is a *state record* type. A state record contains arguments, e.g., the size of a hash table to be allocated in the *open* procedure, and *state*, e.g., the location of a hash table. All state information of an iterator is kept in its state record and there are no "static" variables; thus, an algorithm may be used multiple times in a query by including more than one state record in the query.

All manipulation and interpretation of data objects, e.g., comparisons and hashing, is passed to the iterators by means of pointers to the entry points of appropriate *support functions*. Each of these support functions uses an argument allowing interpreted or compiled query evaluation, as described earlier for file scan predicates. Without the support functions, Volcano's iterators are empty *algorithm shells* that cannot perform any useful work. In effect, the split into algorithm shells and support functions separates control and iteration over sets from interpretation of records or objects. This separation is one of the cornerstones of Volcano's data model independence and extensibility which will be discussed Section 4.

Iterators can be nested and then operate similar to coroutines [28]. In queries involving more than one operator (i.e., almost all queries), state records are linked together by means of *input* pointers. The input pointers are also kept in the state records. They are pointers to structures that represent the input- or subplans. Such a structure, called *QEP* structure for query evaluation plan, includes four pointers. These point to the input operator's state record and to the entry points of the three procedures implementing the operator (*open*, *next*, and *close*). Figure 2

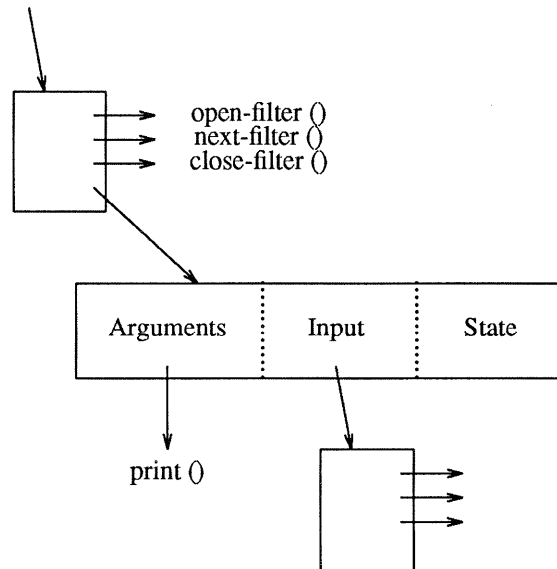


Figure 2. *QEP* Structures and State Record.

shows one operator in a complex tree. Purpose and capabilities of the *filter* operator will be discussed shortly; one of its possible functions is to print items of a stream using a function passed to the filter operator. The *QEP* structure at the top gives access to the functions as well as to the state record. With a pointer to this *QEP* structure, the functions can be called and their local state can be passed to them as a procedure argument. The functions themselves, e.g., *open-filter*, can use the input pointer contained in the state record to invoke the input operator's functions.

Using Volcano's standard form of iterators, an operator does not need to know what kind of operator produces its input, or whether its input comes from a complex query tree or from a simple file scan. We call this concept *anonymous inputs* or *streams*. Streams are a simple but powerful abstraction that allows combining any number and any kind of operators to evaluate a complex query, a second cornerstone to Volcano's extensibility. Together with the iterator control paradigm, streams represent the most efficient execution model in terms of time (overhead for synchronizing operators) and space (number of records that must reside in memory concurrently) for single-process query evaluation.

Calling *open* for the top-most operator results in instantiations for the associated state record's state, e.g., allocation of a hash table, and in *open* calls for all inputs. In this way, all iterators in a query are initiated recursively.

In order to process the query, *next* for the top-most operator is called repeatedly until it fails with an *end-of-stream* indicator. The top-most operator calls the *next* procedure of its input if it needs more input data to produce an output record. Finally, the *close* call recursively "shuts down" all iterators in the query. This model of query execution matches very closely the one being included in the E database implementation language [98] and the query executor of the Starburst relational database system [60].

A number of query and environment parameters may influence policy decisions during *opening* a query evaluation plan, e.g., query predicate constants and system load information. Such parameters are passed between all *open* procedures in Volcano with a parameter called *bindings*. This is a typeless pointer that can be used to pass information for policy decisions. Such policy decisions are implemented using support functions again. For example, the module implementing hash join allows dynamic determination of the size of a hash table — another example of the separation of mechanism and policy. This *bindings* parameter is particularly useful in dynamic query evaluation plans, which will be discussed later in Section 5.

The tree-structured query evaluation plan is used to execute queries by demand-driven dataflow. The return value of a *next* operation is, besides an status indicator, a structure called *NEXT_RECORD* which consists of a RID and a record address in the buffer pool. This record is pinned in the buffer. The protocol about fixing and unfixing records is as follows. Each record pinned in the buffer is *owned* by exactly one operator at any point in time. After receiving a record, the operator can hold on to it for a while, e.g., in a hash table, unfix it, e.g., when a predicate fails, or pass it on to the next operator. Complex operations that create new records, e.g., join, have to fix their output records in the buffer before passing them on, and have to unfix their input records. Since this could result in a large number of buffer calls (one per record in each operator in the query), the interface to the buffer manager was recently redesigned such that it will require a total of two buffer calls per cluster on the producer side (e.g., a file scan) independently of how many records a cluster contains, and only one buffer call per cluster on the consumer side [55].

A *NEXT_RECORD* structure can point to one record only. All currently implemented query processing algorithms pass complete records between operators, e.g., join creates new, complete records by copying fields from two input records. It can be argued that creating complete new records and passing them between operators is prohibitively expensive. An alternative is to leave original records in the buffer as they were retrieved from the stored data, and compose *NEXT_REC* pairs, triples, etc., as intermediate results. Although this alternative results in less

memory-to-memory copying, it is not implemented explicitly because Volcano already provides the necessary mechanisms, namely the *filter* iterator (see next subsection) that can replace each record in a stream by a RID-pointer pair or vice versa.

In summary, demand-driven dataflow is implemented by encoding operators as iterators, i.e., with *open*, *next*, and *close* procedures, since this scheme promises generality, extensibility, efficiency, and low overhead. The next few subsections describe some of Volcano's existing iterators in more detail. In very few modules, the described operators provide much of the functionality of other query evaluation systems through generality and separation of mechanisms and policies. Furthermore, the separation of set processing control (iteration) from item interpretation and manipulation provides this functionality independently from any data model.

3.2.1. Scans, Functional Join, and Filter

The first scan interface was discussed with the file system. The second interface to scans, both file scans and B⁺-tree scans, provides an iterator interface suitable for query processing. The *open* procedures open the file or B⁺-tree and initiate a scan using the scan procedures of the file system level. The file name or closed file descriptor are given in the state record as are an optional predicate and bounds for B⁺-tree scans. Thus, the two scan interfaces are functionally equivalent. Their difference is that the file system scan interface is used by various internal modules, e.g., by the device module for the device table of contents, while the iterator interface is used to provide leaf operators for query evaluation plans.

Typically, B⁺-tree indices hold keys and RIDs in their leaves. In order to use B⁺-tree indices, the records in the data file must be retrieved. In Volcano, this look-up operation is split from the B⁺-tree scan iterator and is performed by the *functional join* operator. This operator requires a stream of records containing RIDs as input and either outputs the records retrieved using the RIDs or it composes new records from the input records and the retrieved records, thus "joining" the B⁺-tree entries and their corresponding data records [74].

B⁺-tree scan and functional join are separated for a number of reasons. First, it is not clear that storing data in B⁺-tree leaves never is a good idea. At times, it may be desirable to experiment with having other types of information associated with look-up keys, e.g., primary keys in secondary indices in ARBRE [78]. The original Ingres design [114], for example, allowed B⁺-tree as primary file organization, but this was abandoned in the commercial version of Ingres because secondary indices must associate primary keys (not RIDs) with secondary index keys, and the lookup procedure for a secondary key must effectively traverse two B⁺-trees to retrieve a record. Second, this

separation allows experimentation with manipulation of RID-lists for complex queries. While Blasgen and Eswaran [13] found nested loops join and merge join always preferable over RID-list joins, their conclusion may not hold for other operations and database systems, e.g., bibliographic search and retrieval systems, or for more recent system environments with large memories [106]. Third, while functional join is currently implemented rather naively, this operation can be made more intelligent. If the functional join operator considers a window of RIDs at a time instead of only a single RID, it may be possible to avoid some buffer faults and to save on disk seek operations if the records are retrieved according to physical location. Fourth, we have generalized functional join to intelligently assemble complex objects consisting of many records in multiple nesting levels. A number of object identifiers (OIDs) or RIDs will be used, selectively based on field values of the input record or of retrieved component objects [49, 54]. In summary, separating index search and record retrieval is another example for providing mechanisms in Volcano to allow for experiments with policies, a design principle employed to ensure that the Volcano software would be flexible and extensible.


In general, selection predicates are applied by the file or B⁺-tree scan operators. Selections that cannot be applied within scans, are performed by the *filter* operator. For example, bit vector filtering [3] can be realized as a selection using the filter operator. The filter operator can actually perform three functions, depending on presence or absence of corresponding support functions in the state record. The *predicate* function applies a selection predicate. The *transform* function creates a new record, typically of a new type, from each old record. An example would be a relational projection (without duplicate elimination). More complex examples include compression and decompression, reducing a stream of records to RID-pointer pairs, changes in codes and representations, and arithmetic. Finally, the *apply* function is invoked once on each record for the benefit of its side effects. Typical examples are updates and printing⁴. Another example is creating a filter for bit vector filtering. In other words, the filter operator is a very versatile single-input single-output operator that can be used for a variety of purposes. Bit vector filtering is an example for a special version of separation of policy and mechanism, namely the rule not to provide an operation that can be composed easily and efficiently using two or more existing operations.

⁴ Notice that updates are done within streams and query evaluation plans. Thus, Volcano plans are not only retrieval but also *update* plans. The filter operator is also called the "side-effect operator."

3.2.2. One-to-One Match

Together with the filter operator, the *one-to-one match* operator will probably be among the most frequently used query processing operators in Volcano as it implements a variety of set-matching functions. In a single operator, it realizes join, semi-join, outer join, anti-join, intersection, union, difference, anti-difference, aggregation, and duplicate elimination. The one-to-one match operator is a physical operator like sort, i.e., part of the executable algebra, not a logical operator like the operators of relational algebra. It is the operator that implements all operations in which an item is included in the output depending on the result of a comparisons between a pair of items⁵.

Figure 3 shows the basic principle underlying the one-to-one match operator for binary operations, namely separation the matching and non-matching components of two sets, called *R* and *S* in the Figure 3, and producing appropriate subsets, possibly after some transformation and combination as in the case of a join. Since all these operations require basically the same steps, it was logical to implement them in one general and efficient module. The main difference between the unary and binary operations, e.g., aggregate functions and equi-join, is that the former require comparing items of the same input while the latter require comparing items of two different inputs.



Output	Full Match	Attribute match
A	Difference	Anti-semi-join
B	Intersection	Join, semi-join
C	Difference	Anti-semi-join
A, B		Left outer join
A, C	Anti-difference	Anti-join
B, C		Right outer join
A, B, C	Union	Symmetric outer join

Figure 3. Binary One-to-One Matching.

⁵ Other kinds of restrictive operators are selection (an item is part of the output depending on its own properties and values) relational division (an item is included in the output depending on a set of matches with the items of the divisor relation), and projection / transformation.

Since the implementation of Volcano’s one-to-one match is data-model-independent and all operations on data items are imported via support functions, the module is not restricted to the relational model but can perform set matching functions for arbitrary data types. Furthermore, the hash-based version provides recursive hash table overflow avoidance [39, 72] and resolution similar to hybrid hash join [29, 30, 106] and can therefore handle very large input sizes. The sort-based version of one-to-one match is based on an external sort operator and can also operate on arbitrarily sized inputs.

While there seems to be an abundance of join algorithms, e.g. [4, 12, 13, 16, 29, 30, 35, 39, 43, 52, 70, 71, 74, 80, 82, 84, 85, 87-89, 95-97, 106-108, 110, 112, 124-126, 128, 129], our design goals of extensibility and limited system size led to the choice of only two algorithms (at the current time) to be implemented in Volcano, namely the conventional merge join [12, 13] and hybrid hash join [29, 30, 106]. This choice will also allow experimental research into the duality and tradeoffs between sort- and hash-based query processing algorithms.

The *classic hash join* algorithm (which is the in-memory component of hybrid hash join, as described in [29, 30, 106]) proceeds in two phases. In the first phase, a hash table is built from one input; it is therefore called the *build phase*. In the second phase, the hash table is probed using tuples from the other input to determine matches and to compose output tuples; it is called the *probe phase*. After the probe phase, the hash table and its entries are discarded. Instead, our one-to-one match operator uses a third phase called the *flush phase*, which is needed for aggregate functions and some other operations.

Since the one-to-one match operator is an iterator like all Volcano operators, the three phases are assigned to the *open*, *next*, and *close* functions. *Open* includes the build phase, while the other two phases are included in the *next* function. Successive invocations of the *next* function automatically switch from the probe phase to the flush phase when the second input is exhausted. Closing the build and probe inputs is delayed until one-to-one match is closed because the file system does not allow closing a file on a virtual device until all its records are unpinned. The state to be saved between invocations of *next* includes indicators for probe or flush phase, the current build record, and the current probe record. The current build record is required to locate further build records in the same bucket chain which must also be matched against the current probe record.

If the one-to-one match operator is used to implement relational join, it proceeds exactly like the classical hash join. The build phase requests input tuples from the *build input* and builds hash bucket chains. The probe phase determines matches by looping over members of a bucket and produces output tuples where appropriate. The

flush phase unfixes the tuples in the hash table in the buffer and deallocates the hash table without producing any output. Support functions are used to calculate an item's hash value, to compare two items, and to compose an output item from two input items.

A semi-join is a join in which items of one input are selected depending on matches with items of the other input. In order to implement semi-join efficiently, the one-to-one match operator preserves rather than copies items of the first input. When performing a semi-join that preserves the probe input, a probe input tuple is matched only until a match is found. The probe tuple is then passed on as output. Thus, memory-to-memory copying is avoided. The module performs a semi-join instead of a join if the support function to compose output tuples is omitted.

The opposite semi-join (which preserves build input tuples) requires a more sophisticated implementation. In our implementation, a bit is initialized to *FALSE* for each tuple in the hash table. When a match is found, this bit is set to *TRUE*. The flush phase scans through all buckets and outputs the tuples for which a match was found.

Finally, all three forms of outer join as well as anti-join can be implemented using Volcano's one-to-one match module. Two additional support functions compose result tuples from build or probe input tuples respectively. Input tuples that do not participate in the natural join (and hence have to be augmented with *NULLs*) are identified using the same mechanisms used for semi-joins.

The build phase can be used to eliminate duplicates or to perform an aggregate function in the build input. The one-to-one match module does not require a probe input; if only an aggregation is required without subsequent join, the absence of the probe input in the state record signals to the module that the probe phase should be skipped. For aggregation, instead of inserting a new tuple into the hash table as in the classic hash join, an input tuple is first matched with the tuples in its prospective hash bucket. If a match is found, the new tuple is discarded or its values are aggregated into the existing tuple.

It is possible to improve the operator's dataflow behavior for duplicate elimination by producing output items immediately when encountered. However, a copy would have to be held back to enable duplicate detection. We felt that permitting operators further up in the tree to proceed while keeping the records pinned in the buffer for duplicate detection would increase buffer contention without significant benefits. Therefore, Volcano does not include this variant.

While hash tables in main memory are usually quite fast, a severe problem occurs if the build input does not fit in main memory⁶. This situation is called *hash table overflow*. There are two ways to deal with hash table overflow. First, if a query optimizer is used and can anticipate overflow, it can be avoided by partitioning the input(s). This *overflow avoidance* technique is the basis for the hash join algorithm used in the Grace database machine [39]. Second, overflow files can be used to resolve the problem after it occurs. Several *overflow resolution* schemes have been designed and compared for relational join [29, 30, 33, 40, 106]. Among those, hybrid hash has been considered the most efficient one as it requires the least I/O to temporary files.

For Volcano's one-to-one match, we have adopted hybrid hash join. Compared to the hybrid hash algorithm used in the GAMMA database machine [31], our overflow resolution scheme has several improvements. Items can be inserted into the hash table without copying, i.e., the hash table points directly to records in the buffer as produced by one-to-one match's build input. If input items are not densely packed, however, the available buffer memory can fill up very quickly. Therefore, Volcano one-to-one match has an argument called the *packing threshold*. When the number of items in the hash table reaches this threshold, items are packed densely into overflow files. However, the clusters (pages) of these overflow files are not yet unfixed in the buffer, i.e., no I/O is performed as yet. Only when the number of items in the hash table reaches a second threshold called *spilling threshold* is the first of the partition files unfixed. The clusters of this file are written to disk and the count of items in the hash table accordingly reduced. When this number reaches the spilling threshold again, the next partition is unfixed, etc. If necessary, partitioning is performed recursively as suggested by DeWitt and Gerber [30], with automatically adjusted packing and spilling thresholds.

The initial packing and spilling thresholds can be set to zero; in that case, Volcano's one-to-one match performs overflow avoidance very similar to the join algorithm used in the Grace database machine. Beyond this parameterization of overflow avoidance and resolution, Volcano's one-to-one match also includes optimizations of cluster size and recursion depth similar to the ones used for sorting [16, 56] and for non-uniform hash value distributions, and it can operate on inputs with variable-length records.

⁶ Notice that if an aggregate function or duplicate elimination is performed on the build input, only the *output* of this operation must fit in main memory. In particular when memory is scarce, aggregating into a new, temporary file pays off since it avoids internal fragmentation in the buffer, i.e., the output records are packed densely into clusters.

The extension of the module described so far to set operations started with the observation that the intersection of two union-compatible relations is the same as the natural join of these relations, and can be best implemented as semi-join. The union is the (double-sided) outer join of union-compatible relations. Difference and anti-difference of two sets can be computed using special settings of the algorithm's bells and whistles. Finally, Cartesian product can be implemented by matching successfully all possible pairs of items from the two inputs.

A second version of one-to-one match is based on sorting. Its two modules are a disk-based merge-sort and the actual merge-join. Merge-join has been generalized similarly to hash-join to support semi-join, outer join, anti-join, and set operations. The sort operator has been implemented in such a way that it uses and provides the iterator interface. Opening the *sort* iterator prepares sorted runs for merging. If the number of runs is larger than the maximal fan-in, runs are merged into larger runs until the remaining runs can be merged in a single step. The final merge is performed on demand by the *next* function. If the entire input fits into the sort buffer, it is kept there until demanded by the *next* function. The *sort* operator also supports aggregation and duplicate elimination. It can perform these operations early, i.e., while writing temporary files [10, 73]. The sort algorithm is described and evaluated in detail in [56].

In summary, Volcano's one-to-one match operators are very powerful parts of Volcano's query execution algebra. By separating the control required to operate on sets of items and the interpretation and manipulation of individual items it can perform a variety of set matching tasks commonly found in database query processing for arbitrary data types and data models. The separation of mechanisms and policies for overflow management supports both overflow avoidance or hybrid hash overflow resolution. Implementing sort- and hash-based algorithms in a comparable fashion will allow meaningful experimental research into the duality and tradeoffs between sort- and hash-based query processing algorithms. The iterator interface guarantees that the one-to-one match operator can easily be combined with other operations, including new iterators yet to be designed.

3.2.3. One-to-Many Match

While the one-to-one match operator includes an item in its output depending on a comparison of two items with one another, the one-to-many match operator compares each item with a number of other items to determine whether a new item is to be produced. A typical example is relational division, the relational algebra operator corresponding to universal quantification in relational calculus. There are two versions of relational division in Volcano. The first version, called *naive division*, is based on sorting. The second version, called *hash-division*, utilizes

two hash tables, one on the divisor and one on the quotient. An exact description of the two algorithms and alternative algorithms based on aggregate functions can be found in [48] along with analytical and experimental performance comparisons and detailed discussions of two partitioning strategies, hash-table overflow, and multi-processor implementations. We are currently studying how to generalize these algorithms in a way comparable with the generalizations of aggregation and join, e.g., for a *majority* function [23, 131].

4. Extensibility

A number of database research efforts strive for extensibility, e.g., EXODUS, GENESIS, Postgres, Starburst, DASDBS [93, 104], Cactis [64], and others. Volcano is a very open query evaluation architecture that provides easy extensibility. Let us consider a number of frequently proposed database extensions and how they can be accommodated in Volcano.

First, when extending the object type system, e.g., with a new abstract data type (ADT) like *date* or *box*, the Volcano software is not affected at all because it does not provide a type system for objects. All manipulation of and calculation based on individual objects is performed by support functions. In some regards, Volcano is incomplete (it is not a database system), but by separating set processing and instance interpretation and providing a well-defined interface between them, Volcano is inherently extensible on the level of instance types and semantics.

As a rule, data items that are transferred between operators using some *next* iterator procedure are records. For an extensible or object-oriented database system, this would be an unacceptable problem and limitation. The solution to be used in Volcano is to pass only the root component (record) between operators after loading and fixing necessary component records in the buffer and suitably swizzling inter-record pointers [86]. Very simple objects can be assembled in Volcano with the functional join operator. Generalizations of this operator are necessary for object-oriented or non-first-normal-form database systems, but can be included in Volcano without difficulty. In fact, a prototype for such an *assembly* operator is currently being evaluated [49] for use in the REVELATION object-oriented database systems project [47, 54].

Second, in order to add new functions on individual objects or aggregate functions, e.g., geometric mean, to the database and query processing system, the appropriate support function is required and passed to a query processing routine. In other words, the query processing routines are not affected by the semantics of the support functions as long as interface and return values are correct. The reason why Volcano software is not affected by extensions of the functionality on individual objects is that Volcano's software only provides abstractions and

implementations for dealing with and sequencing *sets* of objects using streams, whereas the capabilities for interpreting and manipulating individual objects are *imported* in the form of support functions.

Third, in order to incorporate a new access method, e.g., multi-dimensional indices in form of R-trees [59], appropriate iterators have to be defined. Notice that it makes sense to perform not only retrieval but also maintenance of storage structures in the form of iterators. For example, if a set of items defined via a predicate (selection) needs to be updated, the iterator or query tree implementing the selection can "feed" its data into a maintenance iterator. The items fed into the maintenance operator should include a reference to the part of the storage structure to be updated, e.g., a RID or a key, and appropriate new values if they have been computed in the selection, e.g., new salaries from old salaries. Updating multiple structures (multiple indices) can be organized and executed very efficiently using nested iterators, i.e., a query evaluation plan. Furthermore, if ordering makes maintenance more efficient as for B-trees [111], an ordering or sort iterator can easily be included in the plan. In other words, it makes sense to think of plans not only as query plans used in retrieval but also as "update plans" or combinations of retrieval and update plans. The stream concept is very open; in particular, anonymous inputs shield existing query processing modules and the new iterators from one another.

Fourth, to include a new query processing algorithm in Volcano, e.g., an algorithm for transitive closure or *nest* and *unnest* operations for nested relations, the algorithm needs to be coded in the iterator paradigm. In other words, the algorithm implementation must provide *open*, *next*, and *close* procedures, and must use these procedures for its input stream or streams. After an algorithm has been brought into this form, its integration with Volcano is trivial. In fact, as the Volcano query processing software became more complex and complete, this was done a number of times.

Extensibility can also be considered in a different context. In the long run, it clearly is desirable to provide an interactive front-end to make using Volcano easier. We are currently working on two front-ends, a non-optimized command interpreter based on Volcano's executable algebra and an optimized one based on a logical algebra or calculus language including query optimization implemented with a new optimizer generator. The translation between plans as produced by an optimizer and Volcano will be accomplished using a module that "walks" query evaluation plans produced by the optimizer and Volcano plans, i.e., state records, support functions, etc. We will also use the optimizing front-end as a vehicle for experimentation with *dynamic query evaluation plans* which are outlined in the next section.

In summary, since Volcano is very modular in its design, extensibility is provided naturally. It could be argued that this is the case only because Volcano does not address the hard problems in extensibility. However, this argument does not hold. Rather, Volcano is only one component of a database system, namely the query execution engine. Therefore, it addresses only a subset of the extensibility problems and ignores a different subset. As a query processing engine, it provides extensibility of its set of query processing algorithms, and it does so in a way that matches well with the extensibility as provided by query optimizer generators. It does not provide other database services and abstractions like a type system and type checking for the support functions since it is not an extensible database system. The Volcano routines assume that query evaluation plans and their support functions are correct. Their correctness has to be ensured before Volcano is invoked, which is entirely consistent with the general database systems concept to ensure correctness at the highest possible level, i.e., as soon as possible after a user query is parsed. The significance of Volcano as an extensible query evaluation system is that it provides a simple but very useful and powerful set of mechanisms for efficient query processing and that it can and has been used as a flexible research tool. Its power comes not only from the fact that it has been implemented following a few, consistent design principles but also from its two *meta-operators* that will be described in the next two sections.

5. Dynamic Query Evaluation Plans

In most database systems, a query embedded in a program written in a conventional programming language is optimized when the program is compiled. The query optimizer must make assumptions about the values of the program variables that appear as constants in the query and the data in the database. These assumptions include that the query can be optimized realistically using guessed "typical" values for the program variables and that the database will not change significantly between query optimization and query evaluation. The optimizer must also anticipate the resources that can be committed to query evaluation, e.g., the size of the buffer or the number of processors. The optimality of the resulting query evaluation plan depends on the validity of these assumptions. If a query evaluation plan is used repeatedly over an extended period of time, it is important to determine when reoptimization is necessary. We are working on a scheme in which reoptimization can be avoided by using a new technique called *dynamic query evaluation plans* [50].

Volcano includes a *choose-plan* operator that allows realization of both multi-plan access modules and dynamic plans. In some sense, it is not an operator as it does not perform any data manipulations. Since it provides control for query execution it is a *meta-operator*. This operator provides the same *open-next-close* protocol as the

other operators and can therefore be inserted into a query plan at any location. The *open* operation decides which of several equivalent query plans to use and invokes the *open* operation for this input. *Open* calls upon a support function for this policy decision, passing it the *bindings* parameter described above. The *next* and *close* operations simply call the appropriate operation for the input chosen during *open*.

Figure 4 shows a very simple dynamic plan. Imagine a selection predicate controlled by a program variable. The index scan and functional join can be much faster than the file scan, but not when the index is non-clustering and a large number of items must be retrieved [81]. Using the plan of Figure 4, however, the optimizer can prepare effectively for both cases, and the application program using this dynamic plan will perform well for any predicate value.

The *choose-plan* operator allows considerable flexibility. If only one *choose-plan* operator is used as top of a query evaluation plan, it implements a multi-plan access module. If multiple *choose-plan* operators are included in a plan, they implement a dynamic query evaluation plan. Thus, all forms of dynamic plans identified in [50] can be realized with one simple and effective mechanism. Note that the *choose-plan* operator does not make the policy decision which of several plans to execute; it only provides the mechanism. The policy is imported using yet another support function. Thus, the decision can be made depending on bindings for query variables (e.g., program variables used as constants in a query predicate), on the resource and contention situation (e.g., the availability of processors and memory), other considerations such as user priority, or all of the above.

The *choose-plan* operator provides significant new freedom in query optimization and evaluation with an extremely small amount of code. Since it is compatible with the query processing paradigm, its presence does not affect the other operators at all, and it can be used in a very flexible way. The operator is another example for Volcano's design principle to provide mechanisms to implement a multitude of policies. We used the same




Figure 4. A Dynamic Query Evaluation Plan.

philosophy when designing and implementing a scheme for parallel query evaluation.

6. Multi-Processor Query Evaluation

A large number of research and development projects have shown over the last decade that query processing in relational database systems can benefit significantly from parallel algorithms, e.g. [1, 4, 7-9, 11, 17, 31-34, 37, 39, 51, 53, 62, 65, 66, 75, 78, 79, 83, 85, 87, 94, 97, 103, 113, 117, 120, 121, 132, 133]. The main reasons why parallelism is relatively easy to exploit in relational query processing systems are that (1) query processing is performed using a tree of operators which can be executed in separate processes and processors connected with pipelines (inter-operator parallelism), and (2) each operator consumes and produces sets which can be partitioned or fragmented into disjoint subsets to be processed in parallel (intra-operator parallelism). Fortunately, the reasons why parallelism is easy to exploit in relational systems does not require the relational data model *per se*, only that queries be processed as sets of data items in a tree of operators. These are exactly the assumptions made in the design of Volcano, and it was therefore logical to parallelize extensible query processing in Volcano.

When Volcano was ported to a multi-processor machine, it was desirable to use all single-process query processing code existing at that point *without any change*. The result is very clean, self-scheduling parallel processing. We call this novel approach the *operator model* of parallelizing a query evaluation engine [53]. In this model, all parallelism issues are localized in one operator that uses and provides the standard iterator interface to the operators above and below in a query tree.

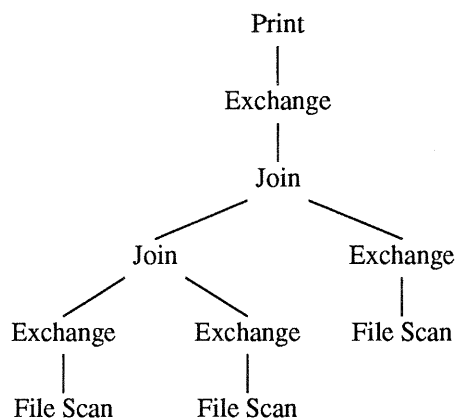


Figure 5a. Operator Model of Parallelization.

The module responsible for parallel execution and synchronization is called the *exchange* iterator in Volcano. Notice that it is an iterator with *open*, *next*, and *close* procedures; therefore, it can be inserted at any one place or at multiple places in a complex query tree. Figure 5a shows a complex query execution plan that includes data processing operators, i.e., file scans and joins, and exchange operators. Figures 5b and 5c will show the processes created when this plan is executed.

This section describes how the *exchange* iterator implements vertical and horizontal parallelism followed by discussions of alternative modes of operation of Volcano's *exchange* operator and modifications to the file system required for multi-process query evaluation. The description goes into a fair amount of detail since the *exchange* operator adds significantly to the power of Volcano. In fact, it represents a new concept in parallel query execution that is likely to prove useful in parallelizing both existing commercial database products and extensible single-process systems. It is described here for shared-memory systems only; considerations for the distributed-memory version are outlined as future work in the last section of this article.

6.1. Vertical Parallelism

The first function of exchange is to provide *vertical parallelism* or pipelining between processes. The *open* procedure creates a new process after creating a data structure in shared memory called a *port* for synchronization and data exchange. The child process is an exact duplicate of the parent process. The exchange operator then takes different paths in the parent and child processes.

The parent process serves as the *consumer* and the child process as the *producer* in Volcano. The exchange operator in the consumer process acts as a normal iterator, the only difference from other iterators is that it receives its input via inter-process communication rather than iterator (procedure) calls. After creating the child process, *open_exchange* in the consumer is done. *Next_exchange* waits for data to arrive via the port and returns them a record at a time. *Close_exchange* informs the producer that it can close, waits for an acknowledgement, and returns.

Figure 5b shows the processes created for vertical parallelism or pipelining by the exchange operators in the query plan of Figure 5a. The exchange operators have created the processes, and are executing on both sides of the process boundaries, hiding the existence of process boundaries from the "work" operators. The fact that the join operators are executing within the same process, i.e., the placement of the exchange operators in the query tree, was arbitrary. The exchange operator provides only the mechanisms for parallel query evaluation, and many other choices (policies) would have been possible. In fact, the mechanisms provided in the operator model tend to be

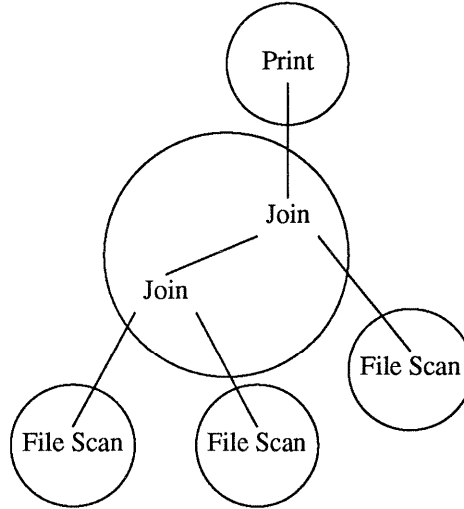


Figure 5b. Vertical Parallelism.

more flexible and amenable to more different policies than in the alternative bracket model [53].

In the producer process, the exchange operator becomes the *driver* for the query tree below the exchange operator using *open*, *next*, and *close* on its input. The output of *next* is collected in *packets*, which are arrays of *NEXT_RECORD* structures. The packet size is an argument in the exchange iterator's state record, and can be set between 1 and 2^{15} records. When a packet is filled, it is inserted into a linked list originating in the *port* and a semaphore is used to inform the consumer about the new packet. Records in packets are fixed in the shared buffer and must be unfixd by a consuming operator.

When its input is exhausted, the exchange operator in the producer process marks the last packet with an *end-of-stream* tag, passes it to the consumer, and waits until the consumer allows closing all open files. This delay is necessary in Volcano because files on virtual devices must not be closed before all their records are unpinned in the buffer. In other words, it is a peculiarity due to other design decisions in Volcano rather than inherent in the exchange iterator or the operator model of parallelization.

The alert reader has noticed that the exchange module uses a different dataflow paradigm than all other operators. While all other modules are based on demand-driven dataflow (iterators, lazy evaluation), the producer-consumer relationship of exchange uses data-driven dataflow (eager evaluation). There are two reasons for this change in paradigms. First, we intend to use the exchange operator also for *horizontal parallelism*, to be described below, which is easier to implement with data-driven dataflow. Second, this scheme removes the need for request

messages. Even though a scheme with request messages, e.g., using a semaphore, would probably perform acceptably on a shared-memory machine, it would create unnecessary control overhead and delays. Since very high degrees of parallelism and very high-performance query evaluation require a closely tied network, e.g., a hypercube, of shared-memory machines, we decided to use a paradigm for data exchange that has been proven to perform well in a "shared-nothing" database machine [31, 34].

A run-time switch of exchange enables *flow control* or *back pressure* using an additional semaphore. If the producer is significantly faster than the consumer, the producer may pin a significant portion of the buffer, thus impeding overall system performance. If flow control is enabled, after a producer has inserted a new packet into the port, it must request the flow control semaphore. After a consumer has removed a packet from the port, it releases the flow control semaphore. The initial value of the flow control semaphore determines how many packets the producers may get ahead of the consumers.

Notice that flow control and demand-driven dataflow are not the same. One significant difference is that flow control allows some "slack" in the synchronization of producer and consumer and therefore truly overlapped execution, while demand-driven dataflow is a rather rigid structure of request and delivery in which the consumer waits while the producer works on its next output. The second significant difference is that data-driven dataflow is easier to combine efficiently with horizontal parallelism and partitioning.

6.2. Horizontal Parallelism

There are two forms of horizontal parallelism which we call *bushy parallelism* and *intra-operator parallelism*. In bushy parallelism, different CPUs execute different subtrees of a complex query tree. Bushy parallelism and vertical parallelism are forms of *inter-operator parallelism*. Intra-operator parallelism means that several CPUs perform the same operator on different subsets of a stored dataset or an intermediate result⁷.

Bushy parallelism can easily be implemented by inserting one or two exchange operators into a query tree. For example, in order to sort two inputs into a merge-join in parallel, the first or both inputs are separated from the merge-join by an exchange operation. The parent process turns to the second sort immediately after forking the

⁷ A fourth form of parallelism is inter-query parallelism, i.e., the ability of a database management system to work on several queries concurrently. In its current version, Volcano does not support inter-query parallelism. A fifth and sixth form of parallelism that can be used for database operations involve hardware vector processing [127] and pipelining in the instruction execution. Since Volcano is a software architecture, and following Boral and DeWitt's analysis in [15], we do not consider hardware parallelism further.

child process that will produce the first input in sorted order. Thus, the two sort operations are working in parallel.

Intra-operator parallelism requires data partitioning. Partitioning of stored datasets is achieved by using multiple files, preferably on different devices [92, 102]. Partitioning of intermediate results is implemented by including multiple queues in a port. If there are multiple consumer processes, each uses its own input queue. The producers use a support function to decide into which of the queues (or actually, into which of the packets being filled by the producer) an output record must go. Using a support function allows implementing round-robin-, key-range-, or hash-partitioning.

Figure 5c shows the processes created for horizontal parallelism or partitioning by the exchange operators in the query plan of Figure 5a. The join operators are executed by three processes while the file scan operators are executed by one or two processes each, typically scanning file partitions on different devices. To obtain this grouping of processes, the only difference to the query plan used for Figure 5b is that the "degree of parallelism" arguments in the exchange state records have to be set to 2 or 3, respectively, and that partitioning support functions must be provided for the exchange operators that transfer file scan output to the join processes. All file scan processes can transfer data to all join processes; however, data transfer between the join operators occurs only within each of the join processes. Unfortunately, this restriction renders this parallelization infeasible if the two joins are on different attributes and partitioning-based parallel join methods are used (rather than a fragment-and-replicate method [4, 38, 112]). For this case, a variant of exchange is supported in Volcano exchange operator




Figure 5c. Horizontal Parallelism.

called *interchange* which is described in the next section.

If an operator or an operator subtree is executed in parallel by a *group* of processes, one of them is designated the *master*. When a query tree is *opened*, only one process is running, which is naturally the master. When a master forks a child process in a producer-consumer relationship, the child process becomes the master within its group. The first action of the master producer is to determine how many slaves are needed by calling an appropriate support function. If the producer operation is to run in parallel, the master producer forks the other producer processes.

After all producer processes are forked, they run without further synchronization among themselves, with two exceptions. First, when accessing a shared data structure, e.g., the port to the consumers or a buffer table, short-term locks must be acquired for the duration of one linked-list insertion. Second, when a producer group is also a consumer group, i.e., there are at least two exchange operators and three process groups involved in a vertical pipeline, the processes that are both consumers and producers synchronize twice. During the (very short) interval between synchronizations, the master of this group creates a port which serves all processes in its group.

When a *close* request is propagated down the tree and reaches the first exchange operator, the master consumer's *close_exchange* procedure informs all producer processes that they are allowed to close down using the semaphore mentioned above in the discussion on vertical parallelism. If the producer processes are also consumers, the master of the process group informs its producers, etc. In this way, all operators are shut down in an orderly fashion, and the entire query evaluation is self-scheduling.

6.3. Variants of the Exchange Operator

There are a number of situations for which the *exchange* operator described so far required some modifications or extensions. In this section, we outline additional capabilities implemented in Volcano's exchange operator. All of these variants have been implemented in the *exchange* operator and are controlled by arguments in the state record.

For some operations, it is desirable to *replicate* or *broadcast* a stream to *all* consumers. For example, one of the two partitioning methods for hash-division [48] requires that the divisor be replicated and used with each partition of the dividend. Another example are fragment-and-replicate parallel join algorithms in which one of the two input relations is not moved at all while the other relation is sent to all processors [4, 38, 112]. To support these algorithms, the exchange operator can be directed to send all records to all consumers, after pinning them appropri-

ately multiple times in the buffer pool. Notice that it is not necessary to copy the records since they reside in a shared buffer pool; it is sufficient to pin them such that each consumer can unpin them as if it were the only process using them.

During implementation and benchmarking of parallel sorting [55, 56], we added two more features to *exchange*. First, we wanted to implement a merge network in which some processors produce sorted streams merge concurrently by other processors. Volcano's *sort* iterator can be used to generate a sorted stream. A *merge* iterator was easily derived from the *sort* module. It uses a single level merge, instead of the cascaded merge of runs used in *sort*. The input of a *merge* iterator is an *exchange*. Differently from other operators, the merge iterator requires to distinguish the input records by their producer. As an example, for a join operation it does not matter where the input records were created, and all inputs can be accumulated in a single input stream. For a merge operation, it is crucial to distinguish the input records by their producer in order to merge multiple sorted streams correctly.

We modified the *exchange* module such that it can keep the input records separated according to their producers. A third argument to *next_exchange* is used to communicate the required producer from the *merge* to the *exchange* iterator. Further modifications included increasing the number of input buffers used by *exchange*, the number of semaphores (including for flow control) used between producer and consumer part of *exchange*, and the logic for *end-of-stream*. All these modifications were implemented in such a way that they support multi-level merge trees, e.g., a parallel binary merge tree as used in [9]. The merging paths are selected automatically such that the load is distributed as evenly as possible in each level.

Second, we implemented a sort algorithm that sorts data randomly partitioned (or "striped" [102]) over multiple disks into a range-partitioned file with sorted partitions, i.e., a sorted file distributed over multiple disks. When using the same number of processors and disks, required two processes per CPU, one to perform the file scan and partition the records and another one to sort them. Creating and running more processes than processors can inflict a significant cost since these processes compete for the CPUs and therefore require operating system scheduling. While the scheduling overhead may not be too significant, in our current environment with a central run queue allowing processes to migrate freely and a large cache associated with each CPU, the frequent cache migration adds a significant cost [123].

In order to make better use of the available processing power, we decided to reduce the number of processes by half, effectively moving to one process per CPU. This required modifications to the exchange operator. Until

then, the exchange operator could "live" only at the top or the bottom of the operator tree in a process. Since the modification, the exchange operator can also be in the middle of a process' operator tree. When the exchange operator is *opened*, it does not fork any processes but establishes a communication port for data exchange. The *next* operation requests records from its input tree, possibly sending them off to other processes in the group, until a record for its own partition is found. This mode of operation was termed *interchange*, and was referred to earlier in the discussion of Figure 5c.

This mode of operation also makes flow control obsolete. A process runs a producer (and produces input for the other processes) only if it does not have input for the consumer. Therefore, if the producers are in danger of overrunning the consumers, none of the producer operators gets scheduled, and the consumers consume the available records.

6.4. File System Modifications

The file system required some modifications to serve several processes concurrently. In order to restrict the extent of such modifications, Volcano currently does not include protection of files and records other than each disk's volume table of contents. Furthermore, typically non-repetitive actions like mounting a device must be invoked by the query root process before or after a query is evaluated by multiple processes.

The most intricate changes were required for the *buffer* module. In fact, making sure the buffer manager would not be a bottleneck in a shared-memory machine was an interesting sub-project independently of database query processing [55]. Concurrency control in the buffer manager was designed to provide a testbed for future research with effective and efficient mechanisms, and not to destroy the separation of policies and mechanisms.

Using one exclusive lock is the simplest way to protect a buffer pool and its internal data structures. However, decreased concurrency would have removed most or all advantages of parallel query processing. Therefore, the buffer uses a two-level scheme. There is a lock for each buffer pool and one for each descriptor (page or cluster resident in the buffer). The buffer pool lock must be held while searching or updating the hash tables and bucket chains. It is never held while doing I/O; thus, it is never held for a long period of time. A descriptor or cluster lock must be held while doing I/O or while updating a descriptor in the buffer, e.g., to decrease its fix count.

Other buffer managers do not use a pool lock but lock each search bucket and the free chain individually, e.g., the buffer manager of Starburst. The advantage is increased concurrency, while the disadvantage is increased

number of locks and lock operations. However, it is much more advantageous to reduce the number of buffer calls than to develop a sophisticated scheme for increased concurrency. Using sixteen 4-MIPS CPUs and two buffer pools, we observed virtually no contention for the pool locks [55].

If a process finds a requested cluster in the buffer, it uses an atomic test-and-lock operation to lock the descriptor. If this operation fails, the pool lock is released, the operation delayed and restarted. It is necessary to restart the buffer operation including the hash table lookup because the process which holds the lock might be replacing the requested cluster. Therefore, the requesting process must wait to determine the outcome of the prior operation. Using this restart-scheme for descriptor locks has the additional benefit of avoiding deadlocks. The four conditions for deadlock are *mutual exclusion*, *hold-and-wait*, *no preemption*, and *circular wait* [63, 67]; Volcano's restart scheme does not satisfy the second condition.

While the locking scheme avoids deadlocks, it does not avoid *convoys* [14]. If a process exhausts its CPU time-slice while holding a "popular" exclusive lock, e.g., on a buffer pool, probably all other processes will block in a convoy until the lock-holding process is re-scheduled and releases the lock. However, since Volcano does not use a "fair" scheduling policy that does not allow reacquiring a lock before all waiting processes held and released the lock, we expect that convoys will quickly evaporate [14]. We intend to investigate the special problem of convoys in shared-memory multi-processors further. A recent investigation of various spin-lock implementations in shared-memory parallel systems showed that the best performance and the least bus traffic was achieved by a queue-lock with a built-in fair scheduler [57]; however, this fair scheduling strategy was exactly the problem identified in [14] that prevented convoys from evaporating. Convoys in shared-memory parallel systems have not been investigated, and we plan on using Volcano for this purpose.

In summary, the exchange module encapsulates parallel query processing in Volcano. It provides a large set of mechanisms useful in parallel query evaluation. Only very few changes had to be made in the buffer manager and the other file system modules to accommodate parallel execution. The most important properties of the exchange module are that it implements three forms of parallel processing within a single module, that it makes parallel query processing entirely self-scheduling, supports a variety of policies, e.g., partitioning schemes or packet sizes, and that it did not require any changes in the existing query processing modules, thus leveraging significantly the time and effort spent on them and allowing easy parallel implementation of new algorithms. It entirely separates data selection, manipulation, derivation, etc. from all parallelism issues, and may therefore prove useful in

parallelizing other systems, both relational commercial and extensible research system.

7. Summary and Conclusions

We have described Volcano, a new query evaluation system that combines compactness, efficiency, extensibility, and parallelism in a dataflow query evaluation system. Compactness is achieved by focusing on few general algorithms. For example, the one-to-one match operator implements join, semi-join, outer join, anti-join, duplication elimination, aggregation, intersection, union, difference, and anti-difference. Extensibility is achieved by implementing only one essential abstraction, streams, and by relying on imported *support functions* for object interpretation and manipulation. The details of streams, e.g., type and structure of their elements, are not part of the stream definition and its implementation, and can be determined at will, making Volcano a *data-model-independent set processor*. The separation of set processing control in *iterators* and object interpretation and manipulation through support functions contributes significantly to Volcano's extensibility.

The Volcano design and implementation was guided by a few simple but generally useful principles. First, Volcano implements *mechanisms* to support *policies* to be determined by a human experimenter or a query optimizer. Second, operators are implemented as iterators to allow efficient transfer of data and control within a single process. Third, a uniform operator interface allows for integration of new query processing operators and algorithms. Fourth, the interpretation and manipulation of stream elements is consistently left open to allow supporting any data model and processing items of any type, shape, and representation. Finally, the encapsulated implementation of parallelism allows developing query processing algorithms in a single-process environment but executing them in parallel. These principles have led to a very flexible, extensible, and powerful query processing engine.

Volcano introduces two novel *meta-operators*. Dynamic query evaluation plans are a new concept introduced in [50] that allow efficient evaluation of queries with free variables. The *choose-plan* meta-operator at the top of a plan or a subplan makes an efficient decision which alternative plan to use when the plan is invoked. Dynamic plans have the potential of increasing the performance of embedded and repetitive queries significantly.

Dataflow techniques are used within processes as well as between processes. Within a process, demand-driven dataflow is implemented by means of streams and iterators. Streams and iterators represent the most efficient execution model in terms of time and space for single-process query evaluation. Between processes, data-driven dataflow is used to pass data between producers and consumers efficiently. If necessary, Volcano's data-driven dataflow can be augmented with flow control or back pressure. Horizontal partitioning can be used both on stored

and intermediate datasets to allow intra-operator parallelism. The design of the *exchange* meta-operator encapsulates the parallel execution mechanism for vertical, bushy, and intra-operator parallelism, and it performs the translations from demand-driven to data-driven dataflow and back [53].

Encapsulating all issues of parallelism control in one operator and thus orthogonalizing data manipulation and parallelism offers an important portability advantage. All data manipulation operators are shielded from parallelism issues, and could be designed, debugged, tuned, and preliminarily evaluated in a single-process environment. To port all Volcano operators to a new parallel machine, only the exchange operator requires appropriate modifications. At the current time, the exchange operator supports parallelism only on shared-memory machines. We are currently working on extending this operator to support query processing on distributed-memory machines while maintaining its encapsulation properties. However, we do not want to give up the advantages of shared memory, namely fast communication and synchronization. A recent investigation demonstrated that shared-memory architectures can deliver near-linear speed-up for limited degrees of parallelism; we observed a speed-up of 14.9 with 16 CPUs for parallel sorting in Volcano [55]. To combine the best of both worlds, we are building our software such that it runs on a closely-tied group, e.g., a hypercube or mesh architecture, of shared-memory parallel machines. Once this version of Volcano's exchange operator and therefore Volcano runs on such machines, we can investigate query processing on hierarchical architectures and heuristics of how CPU and I/O power as well as memory can best be placed and exploited in such machines.

Most of today's parallel machines are built as one of the two the extreme cases of this hierarchical design: a distributed-memory machine uses single-CPU nodes, while a shared-memory machine consists of a single node. Software designed for this hierarchical architecture will run on either conventional design as well as a genuinely hierarchical machine, and will allow exploring tradeoffs in the range of alternatives in between. Thus, the operator model of parallelization also offers the advantage of architecture- and topology-independent parallel query evaluation.

A number of features make Volcano an interesting object for further performance studies. First, the LRU/MRU buffer replacement strategy switched by a keep-or-toss hint needs to be evaluated. Second, using clusters of different sizes on a single device and avoiding buffer shuffling by allocating buffer space dynamically instead of statically require careful evaluation. Third, the duality and tradeoffs between sort- and hash-based query processing algorithms and their implementations will be explored further. Fourth, Volcano allows measuring the

performance of parallel algorithms and identifying bottlenecks on a shared-memory architecture, as demonstrated for instance in [55]. We intend to perform similar studies on distributed-memory and, as they become available, hierarchical architectures. Fifth, the advantages and disadvantages of a separate scheduler process in distributed-memory query processing (as used in GAMMA) will be evaluated. Finally, after data-driven dataflow has been shown to work well on a shared-nothing database machine [31, 34], the combination of demand- and data-driven dataflow should be explored on a network on shared-memory computers.

While Volcano is a working system in its current form, we are considering several extensions and improvements. First, while Volcano currently does very extensive error detection, it does not encapsulate errors in *fail-fast* modules. It would be desirable to modify all modules such that they have all-or-nothing semantics for all requests. This might prove particularly tricky for the exchange module, even more so in a distributed-memory environment. Second, for a more complete performance evaluation, Volcano should be enhanced to a multi-user system that allows inter-query parallelism. Third, to make it a complete data manager and query processor, transactions semantics including recovery should be added.

Volcano is the first operational query evaluation system that combines extensibility and parallelism. We believe that Volcano is a powerful tool for database systems research and education. We are making it available for student use, e.g., for implementation and performance studies, and have given copies to selected outside organizations. We intend to use it in a number of further research projects, including research on the optimization and evaluation of dynamic query evaluation plans [50] and the REVELATION project on query optimization and execution in object-oriented database systems with encapsulated behavior [47, 54].

Acknowledgements

The one-to-one match operators were implemented by Tom Keller starting with existing hash join, hash aggregate, merge join, and sort code. Hash table overflow management was added by Mark Swain. Dynamic query evaluation plans and the *choose-plan* operator were designed and implemented by Karen Ward. We are also very much indebted to the members of the GAMMA and EXODUS projects. Leonard Shapiro contributed to the quality and clarity of the exposition with many insightful comments. David DeWitt, Jim Gray, David Maier, Bill McKenna, Marguerite Murphy, and Mike Stonebraker gave very helpful comments on earlier drafts of this paper. — This research was supported by NSF through grants IRI-8996270 and IRI-8912618, and by the Oregon Advanced Computing Institute (OACIS).

References

1. W. Alexander and G. Copeland, "Process and Dataflow Control in Distributed Data-Intensive Systems", *Proceedings of the ACM SIGMOD Conference*, Chicago, IL., June 1988, 90-98.
2. M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade and V. Watson, "System R: A Relational Approach to Database Management", *ACM Transactions on Database Systems* 1, 2 (June 1976), 97-137.
3. E. Babb, "Implementing a Relational Database by Means of Specialized Hardware", *ACM Transactions on Database Systems* 4, 1 (March 1979), 1-29.
4. C. K. Baru, O. Frieder, D. Kandlur and M. Segal, "Join on a Cube: Analysis, Simulation, and Implementation", *Proceedings of the 5th International Workshop on Database Machines*, 1987.
5. D. S. Batory, "GENESIS: A Project to Develop an Extensible Database Management System", *Proceedings of the Int'l Workshop on Object-Oriented Database Systems*, Pacific Grove, CA., September 1986, 207-208.
6. D. S. Batory, J. R. Barnett, J. F. Garza, K. P. Smith, K. Tsukuda, B. C. Twichell and T. E. Wise, "GENESIS: An Extensible Database Management System", *IEEE Transactions on Software Engineering SE-14*, 11 (November 1988), 1711.
7. R. Bayer, H. Heller and A. Reiser, "Parallelism and Recovery in Database Systems", *ACM Transactions on Database Systems* 5, 2 (June 1980).
8. M. Beck, D. Bitton and W. K. Wilkinson, "Sorting Large Files on a Backend Multiprocessor", *IEEE Transactions on Computers* 37 (1988), 769-778.
9. D. Bitton, H. Boral, D. J. DeWitt and W. K. Wilkinson, "Parallel Algorithms for the Execution of Relational Database Operations", *ACM Transactions on Database Systems* 8, 3 (September 1983), 324-353.
10. D. Bitton and D. J. DeWitt, "Duplicate Record Elimination in Large Data Files", *ACM Transactions on Database Systems* 8, 2 (June 1983), 255-265.
11. D. Bitton Friedland, "Design, Analysis, and Implementation of Parallel External Sorting Algorithms", *Computer Sciences Technical Report 464* (January 1982), University of Wisconsin — Madison.
12. M. Blasgen and K. Eswaran, "On the Evaluation of Queries in a Relational Database System", *IBM Research Report*, San Jose, CA., April 8, 1976.
13. M. Blasgen and K. Eswaran, "Storage and Access in Relational Databases", *IBM Systems Journal* 16, 4 (1977).
14. M. Blasgen, J. N. Gray, M. Mitoma and T. Price, "The Convoy Phenomenon", *Operating Systems Review* 13, 2 (April 1979), 20-25.
15. H. Boral and D. J. DeWitt, "Database Machines: An Idea Whose Time Has Passed? A Critique of the Future of Database Machines", *Proceeding of the International Workshop on Database Machines*, Munich, 1983.
16. K. Bratbergsengen, "Hashing Methods and Relational Algebra Operations", *Proceedings of the Conference on Very Large Data Bases*, Singapore, August 1984, 323-333.
17. K. Bratbergsengen, "Fast Database Systems", *Proc. NATO Supercomputing Conf.*, June 1989.
18. M. J. Carey and D. J. DeWitt, "Extensible Database Systems", *Proceedings of the Islamorada Workshop*, Islamorada, FL., February 1985.
19. M. J. Carey, D. J. DeWitt, J. E. Richardson and E. J. Shekita, "Object and File Management in the EXODUS Extensible Database System", *Proceedings of the Conference on Very Large Data Bases*, Kyoto, Japan, August 1986, 91-100.
20. M. J. Carey, D. J. DeWitt, D. Frank, G. Graefe, J. E. Richardson, E. J. Shekita and M. Muralikrishna, "The Architecture of the EXODUS Extensible DBMS: A Preliminary Report", *Proceedings of the Int'l Workshop on Object-Oriented Database Systems*, Pacific Grove, CA., September 1986, 52-65.
21. M. J. Carey, D. J. DeWitt and S. L. Vandenberg, "A Data Model and Query Language for EXODUS", *Proceedings of the ACM SIGMOD Conference*, Chicago, IL., June 1988, 413-423.

22. M. J. Carey, D. J. DeWitt, G. Graefe, D. M. Haight, J. E. Richardson, D. T. Schuh, E. J. Shekita and S. Vandenberg, "The EXODUS Extensible DBMS Project: An Overview", in *Readings on Object-Oriented Database Systems*, D. M. S. Zdonik (editor), Morgan Kaufman, San Mateo, CA., 1989.
23. J. V. Carlis, "HAS: A Relational Algebra Operator, or Divide Is Not Enough to Conquer", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1986, 254.
24. D. D. Chamberlin, M. M. Astrahan, M. W. Blasgen, J. N. Gray, W. F. King, B. G. Lindsay, R. Lorie, J. W. Mehl, T. G. Price, F. Putzolo, P. G. Selinger, M. Schkolnik, D. R. Slutz, I. L. Traiger, B. W. Wade and R. A. Yost, "A History and Evaluation of System R", *Communications of the ACM* 24, 10 (October 1981), 632-646.
25. H. T. Chou, "Buffer Management of Database Systems", *Ph.D. Thesis*, May 1985.
26. H. T. Chou and D. J. DeWitt, "An Evaluation of Buffer Management Strategies for Relational Database Systems", *Proceedings of the Conference on Very Large Data Bases*, Stockholm, Sweden, August 1985, 127-141.
27. H. T. Chou, D. J. DeWitt, R. H. Katz and A. C. Klug, "Design and Implementation of the Wisconsin Storage System", *Software - Practice and Experience* 15, 10 (October 1985), 943-962.
28. M. Conway, "A Multiprocessor System Design", *Proceedings of the AFIPS Fall Joint Computer Conference*, 1963, 139-146.
29. D. J. DeWitt, R. Katz, F. Olken, L. Shapiro, M. Stonebraker and D. Wood, "Implementation Techniques for Main Memory Database Systems", *Proceedings of the ACM SIGMOD Conference*, Boston, MA., June 1984, 1-8.
30. D. J. DeWitt and R. H. Gerber, "Multiprocessor Hash-Based Join Algorithms", *Proceedings of the Conference on Very Large Data Bases*, Stockholm, Sweden, August 1985, 151-164.
31. D. J. DeWitt, R. H. Gerber, G. Graefe, M. L. Heytens, K. B. Kumar and M. Muralikrishna, "GAMMA - A High Performance Dataflow Database Machine", *Proceedings of the Conference on Very Large Data Bases*, Kyoto, Japan, August 1986, 228-237.
32. D. J. DeWitt, S. Ghandeharizadeh and D. Schneider, "A Performance Analysis of the GAMMA Database Machine", *Proceedings of the ACM SIGMOD Conference*, Chicago, IL., June 1988, 350-360.
33. D. DeWitt and D. Schneider, "A Performance Evaluation of Four Parallel Join Algorithms in a Shared-Nothing Multiprocessor Environment", *Proceedings of the ACM SIGMOD Conference*, Portland, OR, May-June 1989, 110.
34. D. J. DeWitt, S. Ghandeharadizeh, D. Schneider, A. Bricker, H. I. Hsiao and R. Rasmussen, "The Gamma Database Machine Project", *IEEE Transactions on Knowledge and Data Engineering* 2, 1 (March 1990), 44-62.
35. B. C. Desai, "Performance of a Composite Attribute and Join Index", *IEEE Transactions on Software Engineering SE-15*, 2 (February 1989), 142.
36. W. Effelsberg and T. Haerder, "Principles of Database Buffer Management", *ACM Transactions on Database Systems* 9, 4 (December 1984), 560-595.
37. S. Englert, J. Gray, R. Kocher and P. Shah, "A Benchmark of NonStop SQL Release 2 Demonstrating Near-Linear Speedup and Scaleup on Large Databases", *Tandem Computer Systems Technical Report 89.4* (May 1989).
38. R. Epstein and M. Stonebraker, "Analysis of Distributed Data Base Processing Strategies", *Proceedings of the Conference on Very Large Data Bases*, Montreal, Canada, October 1980, 92-101.
39. S. Fushimi, M. Kitsuregawa and H. Tanaka, "An Overview of The System Software of A Parallel Relational Database Machine GRACE", *Proceeding of the Conference on Very Large Data Bases*, Kyoto, Japan, August 1986, 209-219.
40. R. H. Gerber, "The Hash-Partitioned Algorithms", *Preliminary Proposal*, Madison, January 1985.
41. R. Gerber, "Dataflow Query Processing using Multiprocessor Hash-Partitioned Algorithms", *Ph.D. Thesis*, October 1986.

42. R. H. Gerber and D. J. DeWitt, "The Impact of Hardware and Software Alternatives on the Performance of the Gamma Database Machine", *Computer Sciences Technical Report 708* (July 1987), University of Wisconsin — Madison.
43. P. Goyal, H. F. Li, E. Regener and F. Sadri, "Scheduling Page Fetches in Join Operations Using Bc-Trees", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1988, 304-310.
44. G. Graefe, "Rule-Based Query Optimization in Extensible Database Systems", *Ph.D. Thesis*, Madison, August 1987.
45. G. Graefe and D. J. DeWitt, "The EXODUS Optimizer Generator", *Proceedings of the ACM SIGMOD Conference*, San Francisco, CA., May 1987, 160-171.
46. G. Graefe, "Software Modularization with the EXODUS Optimizer Generator", *IEEE Database Engineering*, December 1987.
47. G. Graefe and D. Maier, "Query Optimization in Object-Oriented Database Systems: A Prospectus", in *Advances in Object-Oriented Database Systems*, vol. 334, K. R. Dittrich (editor), Springer-Verlag, September 1988, 358-363.
48. G. Graefe, "Relational Division: Four Algorithms and Their Performance", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA, February 1989, 94-101.
49. G. Graefe, "Set Processing and Complex Object Assembly in Volcano and the REVELATION Project", *Oregon Graduate Center, Computer Science Technical Report*, Beaverton, OR., June 1989.
50. G. Graefe and K. Ward, "Dynamic Query Evaluation Plans", *Proceedings of the ACM SIGMOD Conference*, Portland, OR, May-June 1989, 358.
51. G. Graefe, "Parallelizing the Volcano Database Query Processor", *Digest of Papers, 35th CompCon Conference*, San Francisco, CA., Feb-Mar 1990, 490-493.
52. G. Graefe, "Heap-Filter Merge Join: A New Algorithm for Joining Medium-Size Inputs", *submitted for publication, also CU Boulder Comp. Sci. Tech. Rep. 471*, May 1990.
53. G. Graefe, "Encapsulation of Parallelism in the Volcano Query Processing System", *Proceedings of the ACM SIGMOD Conference*, Atlantic City, NJ., May 1990, 102.
54. G. Graefe, D. Maier, S. Daniels and T. Keller, "A Software Architecture for Efficient Query Processing in Object-Oriented Database Systems with Encapsulated Behavior", *unpublished manuscript*, April 1990.
55. G. Graefe and S. S. Thakkar, "Tuning a Parallel Database Algorithm on a Shared-Memory Multiprocessor", *submitted for publication, also CU Boulder Comp. Sci. Tech. Rep. 470*, April 1990.
56. G. Graefe, "Parallel External Sorting in Volcano", *submitted for publication, also CU Boulder Comp. Sci. Tech. Rep. 459*, February 1990.
57. G. Graunke and S. Thakkar, "Synchronization Algorithms for Shared-Memory Multiprocessors", *IEEE Computer* 23, 6 (June 1990), 60.
58. L. Guibas and R. Sedgwick, "A Dichromatic Framework for Balanced Trees", *Proceedings 19th Symposium on the Foundations of Computer Science*, 1978.
59. A. Guttman, "R-Trees: A Dynamic Index Structure for Spatial Searching", *Proceedings of the ACM SIGMOD Conference*, Boston, MA., June 1984, 47-57.
60. L. M. Haas, W. F. Cody, J. C. Freytag, G. Lapis, B. G. Lindsay, G. M. Lohman, K. Ono and H. Pirahesh, "An Extensible Processor for an Extended Relational Query Language", *Computer Science Research Report*, San Jose, CA., April 1988.
61. L. Haas, W. Chang, G. Lohman, J. McPherson, P. F. Wilms, G. Lapis, B. Lindsay, H. Pirahesh, M. J. Carey and E. Shekita, "Starburst Mid-Flight: As the Dust Clears", *IEEE Transactions on Knowledge and Data Engineering* 2, 1 (March 1990), 143-160.
62. T. Haerder, H. Schoning and A. Sikeler, "Parallel Query Evaluation: A New Approach To Complex Object Processing", *IEEE Database Engineering* 12, 1 (March 1989), 58-64.
63. R. C. Holt, "Some Deadlock Properties of Computer Systems", *ACM Computing Surveys* 4, 3 (September 1972), 179-196.

64. S. E. Hudson and R. King, "Cactis: A Self-Adaptive, Concurrent Implementation of an Object-Oriented Database Management System", *ACM Transaction on Database Systems* 14, 3 (September 1989), 291.
65. H. Itoh, M. Abe, C. Sakama and Y. Mitomo, "Parallel Control Techniques for Dedicated Relational Database Engines", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1987, 208.
66. B. R. Iyer and D. M. Dias, "System Issues in Parallel Sorting for Database Systems", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA, February 1990, 246.
67. E. G. C. Jr., M. J. Elphick and A. Shoshani, "System Deadlocks", *ACM Computing Surveys* 3, 2 (June 1971), 67-78.
68. R. H. Katz, "Performance Enhancement for Relational Systems through Query Compilation", *Proceedings National Computer Conference*, 1979, 741-747.
69. T. Keller and G. Graefe, "The One-to-One Match Operator of the Volcano Query Processing System", *Oregon Graduate Center, Computer Science Technical Report*, Beaverton, OR., June 1989.
70. W. Kim, "A New Way to Compute the Product and Join of Relations", *Proceedings of the ACM SIGMOD Conference*, Santa Monica, CA., May 1980, 179-187.
71. M. Kitsuregawa, L. Harada and M. Takagi, "Join Strategies on KD-Tree Indexed Relations", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA, February 1989, 85.
72. M. Kitsuregawa, et al., "Architecture and performance of relational algebra machine GRACE", *Proc. Int. Conf. Parallel Processing*, 1984, 241-250.
73. A. Klug, "Access Paths in the 'ABE' Statistical Query Facility", *Proceedings of the ACM SIGMOD Conference*, Orlando, FL., June 1982, 161-173.
74. R. P. Kooi, "The Optimization of Queries in Relational Databases", *Ph.D. Thesis*, September 1980.
75. S. C. Kwan, "External Sorting: I/O Analysis and Parallel Processing Techniques", *Ph.D. Thesis*, January 1986.
76. B. W. Lampson, "Hints for Computer System Design", *Proceedings 9th Symp. on Operating System Prin., Operating Systems Review* 17, 5 (October 1983), 33-48.
77. R. A. Lorie and J. F. Nilsson, "An Access Specification Language for a Relational Database Management System", *IBM Journal of Research and Development* 23, 3 (May 1979), 286-298.
78. R. Lorie, J. Daudenarde, G. Hallmark, J. Stamos and H. Young, "Adding Intra-Transaction Parallelism to an Existing DBMS: Early Experience", *IEEE Database Engineering* 12, 1 (March 1989), 58-64.
79. R. A. Lorie and H. C. Young, "A Low Communication Sort Algorithm for a Parallel Database Machine", *IBM Research Report 6669* (February 1989).
80. H. Lu and M. Carey, "Some Experimental Results on Distributed Join Algorithms in a Local Network", *Proceedings of the Conference on Very Large Data Bases*, Stockholm, Sweden, August 1985, 292-304.
81. L. F. Mackert and G. M. Lohman, "Index Scans Using a Finite LRU Buffer: A Validated I/O Model", *ACM Transaction on Database Systems* 14, 3 (September 1989), 401.
82. M. J. Menon and D. K. Hsiao, "Design and Analysis of a Relational Join Operation for VLSI", *Proceeding of the Conference on Very Large Data Bases*, Cannes, France, September 1981, 44-55.
83. J. Menon, "A Study of Sort Algorithms for Multiprocessor Database Machines", *Proceeding of the Conference on Very Large Data Bases*, Kyoto, Japan, August 1986, 197-206.
84. T. H. Merrett, Y. Kambayashi and H. Yasuura, "Scheduling of Page-Fetches in Join Operations", *Proceeding of the Conference on Very Large Data Bases*, Cannes, France, September 1981, 488-498.
85. K. P. Mikkilineni and S. Y. W. Su, "An Evaluation of Relational Join Algorithms in a Pipelined Query Processing Environment", *IEEE Transactions on Software Engineering SE-14*, 6 (June 1988), 838.
86. J. E. B. Moss, "Working with Persistent Objects: To Swizzle or Not to Swizzle", *Computer and Information Science Technical Report 90-38* (May 1990), University of Massachusetts at Amherst.
87. M. C. Murphy and D. Rotem, "Effective Resource Utilization for Multiprocessor Join Execution", *LBL Techical Report 26601* (1989).

88. T. Nakayama, M. Hirakawa and T. Ichikawa, "Architecture and Algorithm for Parallel Execution of a Join Operation", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., April 1984, 160.
89. M. Nakayama, M. Kitsuregawa and M. Takagi, "Hash-Partitioned Join Method Using Dynamic Destaging Strategy", *Proceedings of the Conference on Very Large Databases*, Long Beach, CA., August 1988, 468-478.
90. M. N. Nelson, B. B. Welch and J. K. Ousterhout, "Caching in the Sprite Network File System", *ACM Transactions on Computer Systems* 6, 1 (February 1988), 134-154.
91. J. K. Ousterhout, A. R. Chersonson, F. Douglass, M. N. Nelson and B. B. Welch, "The Sprite Network Operating System", *IEEE Computer* 21, 2 (February 1988), 23-26.
92. D. A. Patterson, G. Gibson and R. H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)", *Proceedings of the ACM SIGMOD Conference*, Chicago, IL., June 1988, 109-116.
93. H. B. Paul, H. J. Schek, M. H. Scholl, G. Weikum and U. Deppisch, "Architecture and Implementation of the Darmstadt Database Kernel System", *Proceedings of the ACM SIGMOD Conference*, San Francisco, CA., May 1987, 196-207.
94. H. Pirahesh, C. Mohan, J. Cheng, T. S. Liu and P. Selinger, "Parallelism in Relational Database Systems: Architectural Issues and Design Approaches", *IBM Research Report*, San Jose, CA, March 1990.
95. G. Z. Qadah and K. B. Irani, "The Join Algorithms on a Shared-Memory Multiprocessor Database Machine", *IEEE Transactions on Software Engineering SE-14*, 11 (November 1988), 1668.
96. G. Z. Qadah, "Filter-Based Join Algorithms on Uniprocessor and Distributed-Memory Multiprocessor Database Machines", *Lecture Notes in Computer Science* 303 (April 1988), 388, Springer Verlag.
97. J. P. Richardson, H. Lu and K. Mikkilineni, "Design and Evaluation of Parallel Pipelined Join Algorithms", *Proceedings of the ACM SIGMOD Conference*, San Francisco, CA., May 1987, 399-409.
98. J. E. Richardson and M. J. Carey, "Programming Constructs for Database System Implementation in EXODUS", *Proceedings of the ACM SIGMOD Conference*, San Francisco, CA., May 1987, 208-219.
99. J. E. Richardson, M. J. Carey, D. J. DeWitt and D. T. Schuh, "Persistence in EXODUS", *Proceedings of A Workshop on Persistent Object Systems: Their Design, Implementation, and Use*, Appin, Scotland, August 1987.
100. J. E. Richardson and M. J. Carey, "Persistence in the E Language: Issues and Implementation", *Computer Sciences Technical Report 791* (September 1988), University of Wisconsin — Madison.
101. G. M. Sacco and M. Schkolnik, "A Mechanism for Managing the Buffer Pool in a Relational Database System Using the Hot Set Model", *Proceeding of the Conference on Very Large Data Bases*, Mexico City, Mexico, September 1982, 257-262.
102. K. Salem and H. Garcia-Molina, "Disk Striping", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1986, 336.
103. B. Salzberg, A. Tsukerman, J. Gray, M. Stewart, S. Uren and B. Vaughan, "FastSort: An Distributed Single-Input Single-Output External Sort", *Proceedings of the ACM SIGMOD Conference*, Atlantic City, NJ., May 1990, 94.
104. H. J. Schek, H. B. Paul, M. H. Scholl and G. Weikum, "The DASDBS Project: Objectives, Experiences, and Future Prospects", *IEEE Transactions on Knowledge and Data Engineering* 2, 1 (March 1990), 25-43.
105. P. Schwarz, W. Chang, J. C. Freytag, G. Lohman, J. McPherson, C. Mohan and H. Pirahesh, "Extensibility in the Starburst Database System", *Proceedings of the Int'l Workshop on Object-Oriented Database Systems*, Pacific Grove, CA., September 1986, 85-92.
106. L. D. Shapiro, "Join Processing in Database Systems with Large Main Memories", *ACM Transactions on Database Systems* 11, 3 (September 1986), 239-264.
107. E. J. Shekita and M. J. Carey, "A Performance Evaluation of Pointer-Based Joins", *Proceedings of the ACM SIGMOD Conference*, Atlantic City, NJ., May 1990, 300.
108. R. Shultz and I. Miller, "Tree Structured Multiple Processor Join Methods", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1987, 190.

109. A. Sikeler, "VAR-PAGE-LRU: A Buffer Replacement Algorithm Supporting Different Page Sizes", *Lecture Notes in Computer Science 303* (April 1988), 336, Springer Verlag.
110. J. M. Smith and P. Y. T. Chang, "Optimizing the Performance of a Relational Algebra Database Interface", *Communications of the ACM 18*, 10 (October 1975), 568-579.
111. J. Srivastava and C. V. Ramamoorthy, "Efficient Algorithms for Maintenance of Large Database Indexes", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1988, 402-408.
112. J. W. Stamos and H. C. Young, "A Symmetric Fragment and Replicate Algorithm for Distributed Joins", *Technical Report RJ7188* (December 5, 1989), IBM Almaden Research Lab.
113. H. S. Stone, "Parallel Querying of Large Databases: A Case Study", *IEEE Computer 21*, 10 (October 1987), 11-21.
114. M. Stonebraker, E. Wong, P. Kreps and G. D. Held, "The Design and Implementation of INGRES", *ACM Transactions on Database Systems 1*, 3 (September 1976), 189-222.
115. M. Stonebraker, "Retrospection on a Database System", *ACM Transactions on Database Systems 5*, 2 (June 1980), 225-240.
116. M. Stonebraker and L. A. Rowe, "The Design of POSTGRES", *Proceedings of the ACM SIGMOD Conference*, Washington, DC., May 1986, 340-355.
117. M. Stonebraker, R. Katz, D. Patterson and J. Ousterhout, "The Design of XPRS", *Proceedings of the Conference on Very Large Databases*, Long Beach, CA., August 1988, 318-330.
118. M. Stonebraker, P. Aoki and M. Seltzer, "Parallelism in XPRS", *UCB/Electronics Research Lab. Memorandum M89/16*, Berkeley, February 1989.
119. M. Stonebraker, L. A. Rowe and M. Hirohama, "The Implementation of Postgres", *IEEE Transactions on Knowledge and Data Engineering 2*, 1 (March 1990), 125-142.
120. S. Y. W. Su and K. P. Mikkilineni, "Parallel Algorithms and Their Implementation in MICRONET", *Proceeding of the Conference on Very Large Data Bases*, Mexico City, Mexico, September 1982, 310-324.
121. S. Y. W. Su, K. P. Mikkilineni, R. A. Liuzzi and Y. C. Chow, "A Distributed Query Processing Strategy Using Decomposition, Pipelining, and Intermediate Result Sharing Techniques", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1986, 94.
122. A. S. Tanenbaum, *Operating Systems – Design and Implementation*, Prentice Hall, Englewood Cliffs, NJ, 1987.
123. S. S. Thakkar and M. Sweiger, "Performance of an OLTP Application on Symmetry Multiprocessor System", *Proc. 17th Annual Int'l Symposium on Computer Architecture, ACM SIGARCH Computer Architecture News 18*, 2 (June 1990), 228.
124. J. A. Thom, K. Ramamohanarao and L. Naish, "A Superjoin Algorithm for Deductive Databases", *Proceeding of the Conference on Very Large Data Bases*, Kyoto, Japan, August 1986, 189-196.
125. S. Todd, "PRTV: An efficient implementation for large relational data bases", *Proceedings of the Conference on Very Large Data Bases*, 1975, 554-556.
126. S. J. P. Todd, "The Peterlee Relational Test Vehicle – A System Overview", *IBM Systems Journal 15*, 4 (1976), 285-308.
127. S. Torii, K. Kojima, Y. Kanada, A. Sakata, S. Yoshizumi and M. Takahashi, "Accelerating Nonnumerical Processing by an Extended Vector Processor", *Proceedings of the IEEE Conference on Data Engineering*, Los Angeles, CA., February 1988, 194-201.
128. P. Valduriez and G. Gardarin, "Join and Semijoin Algorithms for a Multiprocessor Database Machine", *ACM Transactions on Database Systems 9*, 1 (March 1984), 133-161.
129. P. Valduriez, "Join Indices", *ACM Transaction on Database Systems 12*, 2 (June 1987), 218-246.
130. S. L. Vandenberg and D. J. DeWitt, "An Algebra for Complex Objects with Arrays and Identity", *Computer Sciences Technical Report 918* (March 1990), University of Wisconsin — Madison.
131. K. Y. Whang, A. Malhotra, G. Sockut and L. Burns, "Supporting Universal Quantification in a Two-Dimensional Database Query Language", *Proceedings of the IEEE Conference on Data Engineering*, Los

Angelos, CA, February 1990, 68.

132. E. Wong and R. H. Katz, "Distributing a Database for Parallelism", *Proceedings of the ACM SIGMOD Conference*, San Jose, CA., May 1983, 23-29.
133. H. Zeller, "Parallel Query Execution in NonStop SQL", *Digest of Papers, 35th CompCon Conference*, San Francisco, CA., Feb-Mar 1990, 484-487.

Table of Contents

Abstract	1
1. Introduction	1
2. Related Work	3
3. Volcano System Design	5
3.1. The File System	6
3.2. Query Processing	11
3.2.1. Scans, Functional Join, and Filter	15
3.2.2. One-to-One Match	17
3.2.3. One-to-Many Match	21
4. Extensibility	22
5. Dynamic Query Evaluation Plans	24
6. Multi-Processor Query Evaluation	26
6.1. Vertical Parallelism	27
6.2. Horizontal Parallelism	29
6.3. Variants of the Exchange Operator	31
6.4. File System Modifications	33
7. Summary and Conclusions	35
Acknowledgements	37
References	38

List of Figures

Figure 1. Query Preparation	11
Figure 2. QEP Structures and State Record	12
Figure 3. Binary One-to-One Matching	17
Figure 4. A Dynamic Query Evaluation Plan	25
Figure 5a. Operator Model of Parallelization	26
Figure 5b. Vertical Parallelism	27
Figure 5c. Horizontal Parallelism	30